arXiv:1405.0115v5 [math.RA] 29 Jun 2015

KERNELS IN TROPICAL GEOMETRY AND A JORDAN-HÖLDER THEOREM

TAL PERRI, LOUIS ROWEN

Abstract. A correspondence exists between affine tropical varieties and algebraic objects, following the classical Zariski correspondence between irreducible affine varieties and the prime spectrum of the coordinate algebra in affine algebraic geometry. In this context the natural analog of the polynomial ring over a field is the polynomial semiring over a semifield, but one obtains homomorphic images of coordinate algebras via congruences rather than ideals, which complicates the algebraic theory considerably.

In this paper, we pass to the semifield $F(\lambda_1, \dots, \lambda_n)$ of fractions of the polynomial semiring, for which there already exists a well developed theory of **kernels**, also known as **lattice ordered subgroups**; this approach enables us to switch the structural roles of addition and multiplication and makes available much of the extensive theory of chains of homomorphisms of groups. The parallel of the zero set now is the **1-set**. (Idempotent semifields correspond to lattice ordered groups, and the kernels to normal convex l-subgroups.)

These notions are refined in the language of supertropical algebra to ν-**kernels** and 1^ν-**sets**, lending more precision to tropical varieties when viewed as sets of common roots of polynomials. The ν-kernels corresponding to (supertropical) hypersurfaces are the 1^ν-sets of corner internal rational functions. The ν-kernels corresponding to "usual" tropical geometry are the regular, corner-internal ν-kernels.

In analogy to Hilbert's celebrated Nullstellensatz which provides a correspondence between radical ideals and zero sets, we develop a correspondence between 1^ν-sets and a well-studied class of ν-kernels of the rational semifield called **polars**, originating from the theory of lattice-ordered groups. This correspondence becomes simpler and more applicable when restricted to a special kind of ν-kernel, called **principal**, intersected with the ν-kernel generated by F. We utilize this theory to study tropical roots of finite sets of tropical polynomials.

For our main application, we develop algebraic notions such as composition series and convexity degree, along with notions having a geometric interpretation, like reducibility and hyperdimension, leading to a tropical version of the Jordan-Hölder theorem for the relevant class of ν-kernels.

Contents

1. Introduction 1
2. Algebraic Background 5
3. The tropical environment 10
4. Kernels and ν-kernels 13
5. The supertropical connection: Corner loci and K-varieties 17
6. The transition between tropical varieties and K-varieties 21
7. Distinguishing the canonical tropical varieties 24
8. The role of the ν-kernel $\langle F \rangle$ of $\overline{F(\Lambda)}$ in the lattice of regular principal ν-kernels 30
9. Hyperspace-ν-kernels and region ν-kernels 36
10. Polars: an intrinsic description of $\mathcal{K}$-kernels 44
11. Dimension 47
12. The Jordan-Hölder theorem 55
References 56

1. Introduction

This paper is a combination of [29] and [30]. The underlying motive of tropical algebra is that the valuation group of the order valuation (and related valuations) on the Puiseux series field is the ordered group $(\mathbb{R}, +)$ or $(\mathbb{Q}, +)$ (depending on which set one uses for powers in the Puiseux series), which can also be viewed as the max-plus algebra on $(\mathbb{R}, +)$ or $(\mathbb{Q}, +)$. This leads one to a procedure of tropicalization, based on valuations of Puiseux series, which takes us from polynomials over Puiseux series to "tropical" polynomials over the max-plus algebra. One of the main goals of tropical geometry is to study the ensuing varieties. Traditionally, following Zariski, in the affine case, one would pass to the ideal of the polynomial algebra generated by the polynomials defining the variety. However, this does not work well in the tropical world, since one winds up with "too many" varieties for an intuitive theory of dimension. For example, intuitively one would like the

2010 *Mathematics Subject Classification.* Primary: 14T05, 12K10 ; Secondary: 16Y60 .

tropical line to be an irreducible variety of dimension 1. But the intersection of the tropical lines $\lambda_1 + \lambda_2 + 1$ and $\lambda_1 + \lambda_2 + 2$ is the diagonal ray emanating from $(2, 2)$, so one must deal with rays, and the tropical line is the union of three rays. The customary solution to this dilemma is to introduce the "balancing condition," given in rather general terms in Yu [41], but this is tricky in higher dimensions and does not relate directly to the algebraic structure.

Another promising direction, which we follow, is to develop an analog of the celebrated Zariski correspondence, in order to understand tropical geometry in terms of the algebraic structure of various semirings† refining the polynomial algebra over the max-plus algebra. Such a correspondence is one of the traditional ways of treating algebraic geometry in terms of universal algebra, as described in [5, Theorem 1.1] and in considerable generality in [31], and was adapted to tropical algebra by means of corner roots, as described in [24]; a structural description of roots in the "supertropical language" was given in [17].

Let $\mathbb{K} = \mathbb{C}\{\{t\}\}$ be the field of Puiseux series on the variable t, which is the set of formal series of the form $f = \sum_{k=k_0}^{\infty} c_k t^{k/N}$ where $N \in \mathbb{N}$, $k_0 \in \mathbb{Z}$, and $c_k \in \mathbb{C}$. Then $\mathbb{K}$ is an algebraically closed field equipped with the **Puiseux valuation** $val : \mathbb{K}^* = \mathbb{C}\{\{t\}\} \setminus \{0\} \to \mathbb{Q} \subset \mathbb{R}$ defined by

$$val(f) = -\min_{c_k \neq 0}\{k/N\}. \tag{1.1}$$

The **tropicalization** of a Laurent polynomial $f = \sum c_u x^u \in \mathbb{K}[x_1^{\pm 1}, ..., x_n^{\pm 1}]$ is defined to be $trop(f) : \mathbb{R}^{(n)} \to \mathbb{R}$ given by

$$trop(f)(w) = max(val(c_u) + w \cdot u) \tag{1.2}$$

where $u \in \mathbb{N}$ is the power vector and $\cdot$ is the scalar product.

Given a Laurent polynomial $f \in \mathbb{K}[x_1^{\pm 1}, ..., x_n^{\pm 1}]$, one defines its **tropical hypersurface** as

$$T(trop(f)) = \{w \in \mathbb{R}^{(n)} : \text{the maximum in } trop(f) \text{ is achieved at least twice.}\} \tag{1.3}$$

To develop tropical algebraic geometry further, one has the choice either of working directly at the level of polynomials over Puiseux series, and then tropicalizing, or first tropicalizing and then utilizing a Zariski correspondence at the tropical level. The latter is attractive from the point of view of being able to work directly with a simpler concept, and one is led to continue the algebraic study of tropical geometry by means of corner roots of polynomials, developing algebraic geometry over the ordered group $(\mathbb{Q}, +)$; some relevant references are [4, 9, 17, 26]. One tricky issue is that we need to consider intersections of tropical hypersurfaces.

There has also been recent interest in algebraic geometry over monoids, and much of the basic algebraic material can be found in [6]. But many problems arise in formulating algebraic geometry directly over the max-plus algebra, not the least of which is the failure of the max-plus algebra to reflect the uncertainty involved in taking the value of the sum of two polynomials. This was dealt with in [13, 17], by refining the max-plus algebra to a "supertropical semifield†" F and, even more generally, to a "layered" semifield†, i.e., semifield without 0. [9] also treats tropical algebraic geometry in terms of valuations.

When developing tropical affine geometry, it is natural to try to obtain the Zariski-type correspondence between tropical varieties and coordinate semirings. This approach is used for our proposed definition of tropical (affine) variety, in parallel to classical affine algebraic geometry. (Then the usual constructions of schemes and sheaves can be obtained for "semiringed spaces" by noting that negation is not needed in the proofs; the schemes correspond to the congruences defining tropical varieties, and the sheaves to the localizations of the coordinate semirings.) We do not go into the details here, for lack of space.)

Our approach to affine tropical geometry also focuses on a Zariski correspondence parallel to the Zariski correspondence in classical algebraic geometry, which is obtained from the passage between ordered groups and semifields†. Since tropical geometry often is understood in terms of corner roots of polynomials over the target of the order valuation, one should investigate the semiring† $R = F[\lambda_1, \dots, \lambda_n]$ of polynomials over a semifield† F, as well as its homomorphic images.

There is a natural homomorphism ψ from $F[\lambda_1, \dots, \lambda_n]$ to the semiring $\mathrm{Fun}(F^{(n)}, F)$ of functions from $F^{(n)}$ to F, obtained by viewing any polynomial as a function. Since polynomials over the max-plus algebra could have inessential terms which do not affect their values, two different polynomials could have the same image under ψ (such as $5\lambda^2 + \lambda + 7$ and $5\lambda^2 + 7$). But their corner roots must be the same, so it is convenient to work with the image $\psi(F[\lambda_1, \dots, \lambda_n])$, which we denote as $\overline{F[\lambda_1, \dots, \lambda_n]}$. This is particularly apt when one views the coordinate semiring of a variety as the set of polynomial functions restricted to that variety. Viewed in this way, all coordinate semirings are cancellative, and one can define morphisms of varieties in terms of semiring† homomorphisms. (The information lost by passing to $\overline{F[\lambda_1, \dots, \lambda_n]}$ can be accounted for by **region fractions**, cf. Definition 9.2.5. These arise in the decomposition given in Theorem 9.4.8, described presently.)

In general semiring† theory, such homomorphisms are not attained by means of (prime) ideals, but rather via congruences, which are defined as subsets of $R \times R$ rather than of R. Since congruences replace ideals in the study of homomorphic images, one is led to the study of congruences of semirings†. Developing ideas of [1], Joo and Mincheva [22] have developed an elegant theory of prime congruences, to be discussed in Remark 2.4.3.

Several authors, [9, 25, 26] have explored "bend congruences." Even so, the theory of congruences is considerably more complicated than the study of ideals of rings, since it involves substructures of $R \times R$ rather than R. Furthermore, the polynomial algebra over a semifield† need not be Noetherian, cf. [19], although [22] shows that chains of prime congruences are bounded.

The main innovation of this paper, put forth in 2013 by T. Perri in his doctoral dissertation [28], is to switch the roles of multiplication and addition (which is natural enough, since we started with the max-plus algebra). Although mathematically equivalent, this switch enables us to view $\overline{F[\lambda_1, \dots, \lambda_n]}$ as a lattice-ordered (multiplicative) monoid and pass to its group of fractions, focusing on the group structure. This leads us further into the classical theory of lattice-ordered monoids and groups [3, 7, 11, 23, 34, 35] and their corresponding idempotent semifields†.

Put another way, although the polynomial semiring† is not a semifield†, it is a cancellative semiring† when its elements are viewed as functions, so we can view polynomials as functions and then pass to the semifield† of fractions, which is called the semifield† of **rational functions** over F. (The information thereby lost is compensated by Homomorphisms of idempotent semifields† have been studied long ago in the literature [12, 39, 38], where the homomorphic images are described in terms of what they call **(semifield†) kernels**, which as noted in [12] are just the convex (normal) l-subgroups of the corresponding lattice-ordered groups. Since semifield† kernels are subgroups which can be described algebraically, cf. [12, Theorem 3.2], they are more amenable to the classical structure theorems of group theory (Noether isomorphism theorems, Jordan-Hölder theorem, etc.) than congruences. For the supertropical theory, we need the ν-analog of kernels, which we call **ν-kernels**.

The parallel notion to a zero set in algebraic geometry now is the set of points which when substituted into a function give the value 1 instead of 0, so we call these sets **1^ν-sets**. Our ultimate goal is a 1:1 correspondence between sets of corner roots (which in the supertropical language are "ghost roots") of polynomials and 1^ν-sets of rational functions.

A ν-kernel is called **principal** if it is generated by one element. One big advantage of the use of ν-kernels is that the product of finitely many principal ν-kernels is a principal ν-kernel, cf. Theorem 4.1.30, and thus varieties defined as finite intersections of hypersurfaces still can be described in terms of principal ν-kernels. This enables us to study tropical varieties that are not necessarily hypersurfaces.

We come upon various correspondences, but for some reason to be explained below, we need to intersect down to the ν-kernel generated by the constant functions, which we call $\langle F \rangle$, in order to obtain the correspondence given in Corollary 8.0.14:

The lattice of (tangible) simultaneous corner loci with respect to a finite set of polynomials corresponds to the lattice of principal (tangible) ν-kernels of $\langle F \rangle$.

This result, a consequence of Theorem 8.0.13, to be described soon, enables one to transfer the geometric theory to ν-kernels of ν-semifields†. The preparation for its proof requires a careful description of various different kinds of ν-kernels, each with its particular geometric properties, and takes the bulk of this paper.

The application of ν-kernels to tropical geometry, given in [28], is related to the bend congruences of [9], as discussed in Remark 6.1.14.

To obtain our results algebraically, we need a slight modification of the results about semifields† in the literature. As explained in [16] and elaborated in [13, 15], there are other semirings† R covering the usual idempotent max-plus algebra, which provide better tools for examining the algebraic structure arising from valuations on Puiseux series, and these semirings† are no longer embeddable into semifields†. We list the relevant structures in increasing level of refinement:

- The max-plus algebra
- Supertropical algebra
- Layered tropical algebras

In each of the latter three cases, although the multiplicative monoid is not ordered, there is an underlying partially ordered monoid, so we can modify the results about (semifield) ν-kernels in the literature to suit our purposes. Furthermore, in the supertropical and layered situations, there is a "ghost map" ν from R to the set of "ghost elements" $\mathcal{G}$, which enables us to compare elements in R via their images in the ordered monoid $\mathcal{G}$. We treat only the supertropical setting here, since we feel that the theory can be described more concisely in that context. A more thorough description using the layered setting would require even more technical detail.

Let us proceed with a more detailed overview of this paper. Since both ν-semifields† and their ν-kernels may be unfamiliar to many researchers, we review ordered monoids and their semirings† in §2. After an introduction to the supertropical theory in §3, the basic notions of ν-semifield† ν-kernels (and their supertropical analogs) are presented in §4.

The pertinence to tropical geometry is given in §5. Corner roots are replaced by kernel roots, those points whose value at each function in the ν-kernel is ν-congruent to 1; these are called 1^ν-sets, which are the kernel-theoretic analogues of corner loci. The basic Zariski-type correspondence between 1^ν-sets and ν-kernels is given in Theorem 5.3.6. But different ν-kernels could correspond to the same 1^ν-set, cf. Example 8.1.3. This ambiguity is overcome by intersecting with the ν-kernel $\langle F\rangle$. Indeed, one still obtains all the 1^ν-sets, by Proposition 8.1.4, but now Proposition 8.4.2 yields uniqueness in the correspondence, eventually yielding the desired 1:1 correspondence between principal 1^ν-sets and principal ν-subkernels of $\langle \mathscr{R}\rangle$ in Theorem 8.4.3.

The road to Theorem 8.4.3 is somewhat arduous. The interplay of corner roots of polynomials and ν-kernel roots comes in Section 6. In applying the theory to sets of polynomials defining tropical varieties, one quickly sees that there still are too many ν-kernels; for example, some of the varieties defined by ν-kernels do not satisfy basic tropical conditions such as the balancing condition. Example 6.3.4 also displays a certain form of pathology. Thus, there is no way to obtain such basic notions as dimension without specifying which kinds of ν-kernels to yield the tropical varieties arising from hypersurfaces and their intersections. These ν-kernels, called **corner internal** ν-kernels, lie at the heart of our theory, being the ones that provide the transition between ghost roots and kernel roots:

A rational function f is **corner internal** if $f = \frac{h}{g}$ such that every ghost root of $g + h$ is a ν-kernel root of f. A **corner internal** ν-kernel, cf. Definition 7.1.1, is a ν-kernel generated by a corner internal rational function.

There are two main ways of passing from hypersurfaces to corner internal rational function, one given in the basic "hat construction" of §6.1, leading to Theorem 7.1.7, stating that the correspondences $f \mapsto \hat{f}$ and $h \mapsto \underline{h}$ induce 1:1 correspondences between hypersurfaces and 1^ν-sets of corner internal rational functions.

Another, subtler way, of describing corner internal rational functions is given in Definition 7.3.6 and Theorem 7.3.9. In §7.4 we specify the **regular** rational functions, which locally are not the identity, and thus have tropical significance, distinguishing between two general types of nontrivial principal 1^ν-sets in $F^{(n)}$:

- 1^ν-sets not containing a region of dimension n;
- 1^ν-sets that do contain a region of dimension n.

The ν-kernels defining "traditional" tropical hypersurfaces are the regular corner internal ν-kernels.

We further narrow down the ν-kernels of interest, finally showing how one obtains a strong Zariski correspondence using regular principal ν-kernels. There is a basic question as to over which semifields† we should define our ν-kernels. Although most of the results are for ν-kernels of $F(\lambda_1, \dots, \lambda_n)$, the rational fractions of the polynomial semiring† over an arbitrary supertropical semifield† F, in §8.2 we add the condition that the underlying semifield† F be complete. At times we need to assume that F is complete, archimedean, and divisible, in order to obtain principal ν-kernels.

Most of §8 involves intersecting down to $\langle F\rangle$. In Theorem 8.0.13 we formulate the major result:

All (regular) principal ν-kernels of $\langle F\rangle$ are generated by (regular) corner internal principal ν-kernels.

See Example 8.6.2 as to how this result applies to the familiar tropical line. Its proof requires more machinery, and also requires us to take F to be complete under the order topology, say $F = \mathscr{R}$, the supertropical version of the real numbers (viewed as a max-plus structure).

The **wedge decomposition** is a tool introduced in Definition 8.5.2, and in Theorem 8.5.3 it is seen that an intersection of principal ν-subkernels of $\langle \mathscr{R}\rangle$ has a corresponding wedge decomposition.

But we still need more preparation to prove Theorem 8.0.13. Our classes of ν-kernels are further subdivided in §9.4 into "HO-kernels" and another kind of ν-kernel which can be bypassed when restricting from $\overline{F(\Lambda)}$ to $\langle F\rangle$. HO-kernels are decomposed as products of "HS-kernels" and "region ν-kernels" in Theorem 9.4.8:

Every principal ν-kernel $\langle f\rangle$ of $F(\lambda_1, ..., \lambda_n)$ can be written as the intersection of finitely many principal ν-kernels

$$\{K_i : i = 1, ..., q\} \text{ and } \{N_j : j = 1, ..., m\},$$

whereas each K_i is the product of an HS-kernel and a region ν-kernel, and each N_j is a product of "bounded from below" ν-kernels and (complementary) region ν-kernels.

This enables us to reduce to finite decompositions and finally prove Theorem 8.0.13.

As a major application, the decomposition theory we have just described also gives rise to an algebraic description of basic geometric tropical concepts such as dimension, which is seen to match the intuitive

geometric notion. Towards this end, we define a **hyperplane ν-kernel**, or **HP-kernel**, for short, to be a principal ν-kernel of $F(\lambda_1, ..., \lambda_n)$ generated by a single Laurent monomial. More generally, a **hyperspace-fraction ν-kernel**, or **HS-kernel**, is a principal ν-kernel of $F(\lambda_1, ..., \lambda_n)$ generated by a rational function $f \sim_{\langle F \rangle} \sum_{i=1}^{t} |f_i|$ (cf. Definition 8.1.5) where the f_i are non-proportional non-constant Laurent monomials, and corresponds to a hyperplane in polyhedral geometry.

On the other hand, we also encounter **order ν-kernels**. Imposition of an order ν-kernel intersects our variety with a half-space, so does not reduce the dimension. Furthermore, one could have an infinite chain of "parallel" order ν-kernels. For these reasons, our theory of dimension bypasses order ν-kernels and turns to HS-kernels.

The **height** $\mathrm{hgt}(L)$ of an HS-kernel L is given in Definition 11.1.9, and the **convex dimension** then is given in Definition 11.3.10. This dimension is catenary, in the sense of Theorem 11.4.7:

If R is a region ν-kernel and L is an HS-kernel of $\overline{F(\Lambda)}$, then

$$\mathrm{d}_{\mathrm{conv}}(\overline{F(\Lambda)}/LR) = \mathrm{d}_{\mathrm{conv}}(\overline{F(\Lambda)}) - \mathrm{d}_{\mathrm{conv}}(L).$$

In particular, $\mathrm{d}_{\mathrm{conv}}(\overline{F(\Lambda)}/LR) = n - \mathrm{d}_{\mathrm{conv}}(L)$. There is a well-known general argument for Jordan-Hölder type theorems for finite chains in modular lattices, recalled in Theorem 12.0.3, leading us to Theorem 12.0.5:

If a ν-kernel L is irreducible in the lattice of ν-kernels finitely generated by HP-kernels of $\overline{F(\Lambda)}$, then $\mathrm{hgt}(L)$ is its convex dimension. Moreover, every factor of a descending chain of HS-kernels of maximal length is an HP-kernel.

Finally, let us turn to the Nullstellensatz. Several versions for a tropical Hilbert Nullstellensatz have been put forth, including [1, 22]. Our objective here rather is to understand the tropical spirit of the Nullstellensatz in a more general context. Just as the classical Hilbert Nullstellensatz shows that affine varieties correspond to radical ideals of the polynomial algebra, here affine tropical varieties correspond to certain ν-kernels called **polars**, defined intrinsically in terms of the natural orthogonality relation that $f \perp g$ for $f, g \in F(X)$ if their minimum always takes on the value 1.

Polars are best understood as ν-kernels in the completion of the natural image of $F(\lambda_1, \ldots, \lambda_n)$ inside the semiring† of functions from some subset $S \subseteq F^{(n)}$ to F, but we can take the restriction back to $F(\lambda_1, \ldots, \lambda_n)$. Unfortunately, finitely generated polars need not be finitely generated as ν-kernels, which is another reason for us to restrict to bounded ν-kernels.

We discuss some of their basic properties in §10. By Theorem 10.3.7, the polars of $\overline{\mathscr{R}(\Lambda)}$ are precisely those that come from 1^ν-sets, thereby leading up to the fundamental correspondence of Theorem 10.3.9:

There is a $1:1$ correspondence from the polars of $\overline{\mathscr{R}(\Lambda)}$ to the 1^ν-sets in $\mathscr{R}^n$, given by $B \mapsto 1_{\mathrm{loc}}(B)$ and $Z \mapsto \mathcal{K}ern(Z)$, which restricts to a correspondence between the principal polars (as polars) of $\overline{\mathscr{R}(\Lambda)}$ and the principal 1^ν-sets in $\mathscr{R}^n$.

In this way we reduce ν-kernels to finitely generated ν-kernels, which then are principal, and we obtain the appropriate Zariski correspondence in its entirety.

2. Algebraic Background

We start by reviewing some familiar notions that are needed extensively in our exposition.

2.1. Semi-lattice ordered monoids, groups, and semirings.

The passage to the max-plus algebra in tropical mathematics can be viewed via ordered groups and, more generally, ordered monoids and semirings, so we start with them, drawing on the review given in [16]. This material is well known, and largely can be found in [36]. Recall that a monoid $(M, \cdot, 1)$ is just a semigroup with an identity, i.e., unit element, denoted as 1. We work solely with **Abelian** monoids, in which the operation is commutative. Our semigroups without identity are written in additive notation, and monoids are written in multiplicative notation.

Definition 2.1.1. A (set-theoretic) **semi-lattice** L is a partially ordered set with an (associative) binary "sup" operation $\vee$, which means:

$$a, b \leq a \vee b, \quad \text{and} \quad \text{if } a, b \leq c \text{ then } a \vee b \leq c. \tag{2.1}$$

Definition 2.1.2. An element a of a semigroup $(S, +)$ is **idempotent** if $a + a = a$. A **band** is a semigroup in which every element is idempotent. A band $(S, +)$ is **bipotent** if $a + b \in \{a, b\}$ for each pair of elements.

Lemma 2.1.3. *Any semi-lattice can be viewed as a band, where we define $a + b = a \vee b$.*

Proof. To check associativity of addition we have $(a + b) + c = \sup\{a, b, c\} = a + (b + c)$. Furthermore $a + a = a \vee a = a$. □

Definition 2.1.4. A **semi-lattice ordered monoid** is a monoid M that is also a semi-lattice with respect to the operation $\vee$ and satisfies the following property:

$$g(a \vee b) = ga \vee gb \tag{2.2}$$

for all elements $a, b, g \in M$.

A **partially ordered monoid** is a monoid M with a partial order satisfying

$$a \leq b \quad \text{implies} \quad ga \leq gb, \tag{2.3}$$

for all elements $a, b, g \in M$.

An **ordered monoid** is a partially ordered monoid for which the given partial order is a total order.

A **lattice ordered monoid** is a monoid M that is also a lattice with respect to the operations $\vee$ and $\wedge$ satisfying (2.2) for both $\vee$ and $\wedge$.

It follows readily that if $a \leq b$ and $g \leq h$, then $ga \leq gb \leq hb$. Note that the group $(\mathbb{R}, +)$ is ordered, under this definition, but $(\mathbb{R}, \cdot)$ is not.

Every semi-lattice ordered monoid is partially ordered with respect to the partial order given by $a \leq b$ iff $a \vee b = b$.

In case a semi-lattice ordered monoid M is a group $\mathcal{G}$, one defines the **dual semi-lattice** $(\mathcal{G}, \wedge)$ via

$$a \wedge b = (a^{-1} \vee b^{-1})^{-1}. \tag{2.4}$$

Now $\mathcal{G}$ becomes a lattice, and we call $(\mathcal{G}, \vee, \wedge)$ a **lattice ordered group** if $\mathcal{G}$ is a lattice ordered monoid.

Lemma 2.1.5. *Suppose $(\mathcal{G}, \vee)$ is a semi-lattice ordered Abelian group with respect to a given partial order $\leq$. Then $(\mathcal{G}, \wedge)$ is indeed a semi-lattice, with respect to which $\mathcal{G}$ also becomes a lattice ordered group satisfying $(a \vee b)^{-1} = a^{-1} \wedge b^{-1}$.*

Proof. $a \leq b$ iff $b^{-1} = (ab)^{-1}a \leq (ab)^{-1}b = a^{-1}$, so we get the dual partial order $\geq$. Hence $(\mathcal{G}, \wedge)$ is a semi-lattice with respect to $\geq$, since

$$a \wedge b = (a^{-1} \vee b^{-1})^{-1} = \sup\left\{a^{-1}, b^{-1}\right\}^{-1} = \inf\left\{(a^{-1})^{-1}, (b^{-1})^{-1}\right\} = \inf\{a, b\}.$$

Then $g(a \wedge b) = g(a^{-1} \vee b^{-1})^{-1} = (g^{-1}(a^{-1} \vee b^{-1}))^{-1} = (g^{-1}a^{-1} \vee g^{-1}b^{-1})^{-1} = ga \wedge gb$, as desired.

The last assertion follows from (2.4) taking a^{-1}, b^{-1} instead of a, b. □

The duality in Lemma 2.1.5 shows that it suffices to consider only $\vee$ (or only $\wedge$). Every partially ordered group that is a semi-lattice is a semi-lattice ordered group, seen by using g^{-1} instead of g in (2.2).

In any ordered group we define the open interval of distance b around a given by $\{g : ab^{-1} < g < ab\}$, giving rise to the **order topology**.

Although idempotence pervades the theory, it turns out that what is really crucial for many applications is the following well-known fact:

Lemma 2.1.6. *In any band, if $a + b + c = a$, then $a + c = a$.*

Proof. $a = a + b + c = (a + b + c) + c = a + c$. □

Let us call a semigroup **proper** if it satisfies the condition of Lemma 2.1.6.

Remark 2.1.7 ([38, Theorem 4.3]). A proper semigroup cannot have additively invertible elements other than $\mathbf{0}$, since if $a + c = \mathbf{0}$, then $a = a + \mathbf{0} = a + a + c$, implying $a = a + c = \mathbf{0}$.

A submonoid $\mathcal{T}$ of an Abelian monoid M is called **cancellative** if $ab = ac$ for $a \in \mathcal{T}$ and $b, c \in M$ implies $b = c$. In this case, when $\mathcal{T} = M$, we say that the Abelian monoid M is **cancellative**.

A monoid M is **power-cancellative** (called **torsion-free** in [6]) if $a^n = b^n$ for some $n \in \mathbb{N}$ implies $a = b$. Any ordered, cancellative monoid is power-cancellative and infinite.

2.1.1. *Divisibility.*

We say that a monoid M is $\mathbb{N}$-**divisible** (also called **radicalizible** in the tropical literature, but that terminology conflicts with [40] and radical theory) if for each $a \in M$ and $m \in \mathbb{N}$ there is $b \in M$ such that $b^m = a$. For example, $(\mathbb{Q}, +)$ is $\mathbb{N}$-divisible. There is a standard construction of the **divisible hull** of a cancellative monoid M, given in [36], which is semi-lattice ordered when M is semi-lattice ordered; namely,

$$\sqrt[m]{a} \vee \sqrt[n]{b} = \sqrt[mn]{a^n \vee b^m}. \tag{2.5}$$

This also was discussed in [13, Remark 2.3].

2.2. **Semirings without zero.**

Definition 2.2.1. A **semiring**† (semiring without zero, at times called a hemiring) is a set $R := (R, +, \cdot, 1)$ equipped with binary operations $+$ and $\cdot$ and distinguished element $\mathbf{1}_R$ such that:

(i) $(R, +)$ is an Abelian semigroup;
(ii) $(R, \cdot\, , \mathbf{1}_R)$ is a monoid with identity element $\mathbf{1}_R$;
(iii) Multiplication distributes over addition.
(iv) R contains elements r_0 and r_1 with $r_0 + r_1 = \mathbf{1}_R$.

Note that (iv) is automatic if $(R, +)$ is a band, since then $\mathbf{1}_R + \mathbf{1}_R = \mathbf{1}_R$. For the purposes of this paper, a **domain**† is a commutative semiring† whose multiplicative monoid is cancellative.

Definition 2.2.2. A **semifield**† is a domain† in which every element is (multiplicatively) invertible.

(In other words, its multiplicative monoid is an Abelian group. In [38] commutativity is not assumed, but we make this assumption to avoid distraction from our applications.) In particular, the max-plus algebras $(\mathbb{N}, +)$, $(\mathbb{Q}, +)$, and $(\mathbb{R}, +)$ are semifields†, whose multiplication now is given by $+$. We also have the **trivial semifield**† $\{1\}$. In the literature it is customary to write the operations as $\oplus$ and $\odot$, but we use the usual algebraic notation of $+$ and $\cdot$ for addition and multiplication respectively, to emphasize the structural aspects of the theory.

Any semifield† without negatives is proper, by [12, Proposition 20.37].

Digression: The customary definition of semiring [10] also requires the existence of a zero element:

A **semiring** is a semiring† R with a zero element $\mathbf{0}_R$ satisfying

$$a + \mathbf{0}_R = a, \quad a \cdot \mathbf{0}_R = \mathbf{0}_R = \mathbf{0}_R \cdot a, \quad \forall a \in R.$$

(Note that in the definition of semiring† one could then take $r_0 = \mathbf{0}_R$ and $r_1 = \mathbf{1}_R$.) We use semifields† instead of semifields since the zero element usually can be adjoined formally, and often is irrelevant, and the concepts are easier to define when we do not need to treat the zero element separately. Also, the language of semifields† is more appropriate to geometry over tori, which are direct products of groups.

A semiring† R is **idempotent** if the semigroup $(R, +)$ is a band. A semiring† is **bipotent** if the semigroup $(R, +)$ is bipotent. Thus, the max-plus algebra, viewed as a semiring†, is bipotent.

Our basic structures are idempotent semifields†, denoted as F or $\mathbb{S}$ throughout. (Usually F is the underlying semifield†, contained in $\mathbb{S}$.) Let us recall an idea of Green from the theory of idempotent semigroups.

Proposition 2.2.3 ([39, §4])**.** *Any semi-lattice ordered Abelian monoid M becomes an idempotent (commutative) semiring†, which we denote as R, via the usual max-plus procedure; we define the new multiplication on R to be the operation given originally on M, and addition on R is defined as in Lemma 2.1.3:*

$$a + b := a \vee b \tag{2.6}$$

(viewed in M). Conversely, any idempotent commutative semiring† R becomes a semi-lattice ordered Abelian monoid by reading (2.6) *backwards.*

Proof. Distributivity follows from (2.2). Furthermore $a + a = a \vee a = a$.

Conversely, if $a \leq b$ and $b \leq a$ we have $a = a + b = b$, so $\leq$ is antisymmetric; we need to show that $a + b$ is the sup of a and b. In other words, if $a \leq c$ and $b \leq c$ then $a + b \leq c$. But we are given $a + c = c$ and $b + c = c$, so

$$a + b + c = (a + c) + (b + c) = c + c = c,$$

as desired. □

When M is cancellative, then R is a domain†. When M is a group, then R is a semifield†. (This could all be formulated categorically.)

Lemma 2.2.4 ([37, Property 2.6])**.** *Every domain† is torsion-free as a monoid.*

Proof. If $a^n = 1$, then $a(a^{n-1} + a^{n-2} + \cdots + 1) = (1 + a^{n-1} + a^{n-2} + \cdots + a)$, implying $a = 1$. □

Definition 2.2.5. A semiring† R is **ordered** if $(R, +)$ and $(R, \cdot)$ are both ordered monoids.

Remark 2.2.6. The multiplicative group of every ordered semifield† is a torsion-free group, i.e., all of its elements not equal to 1 have infinite order.

Here is another important property:

Definition 2.2.7. A semiring† R is **Frobenius** if it satisfies the identity

$$(a + b)^n = a^n + b^n, \qquad \forall a, b \in R.$$

It is clear that any bipotent semiring† is Frobenius.

Definition 2.2.8. A lattice P is said to be **complete** if all of its bounded subsets have both a supremum and an infimum. A sublattice of P is said to be **completely closed** if all of its subsets have both a supremum and an infimum in P.

An idempotent semiring† is said to be **complete** if its underlying lattice (with $\vee$ as $+$ and $\wedge$ as $\cdot$) is complete.

2.3. **Localization.**

Since we are mainly interested in (proper) idempotent semifields†, we need a method of passing from semirings† to semifields†.

Remark 2.3.1. There is a well-known localization procedure with respect to multiplicative subsets of Abelian monoids, described in detail in [3]. Namely, for any submonoid S of a monoid M, we define an equivalence on $S \times M$ by saying $(s_1, a_1) \sim (s_2, a_2)$ iff there is $s \in S$ such that $ss_2a_1 = ss_1a_2$; then the **localization** $S^{-1}M$ is the set of equivalence classes $\{[(s, a)] : s \in S, a \in M\}$, written as $\frac{a}{s}$. $S^{-1}M$ is a monoid via the operation

$$\frac{a}{s}\frac{a'}{s'} = \frac{aa'}{ss'}.$$

There is a homomorphism $M \mapsto S^{-1}M$ given by $a \mapsto \frac{a}{1}$. This map is 1:1 precisely when $a \neq b$ implies $\frac{a}{1} \neq \frac{b}{1}$; in other words, when the submonoid S of M is cancellative.

If the monoid M itself is cancellative, then localizing with respect to all of M yields its **group of fractions**. In this case $\frac{a_1}{s_1} = \frac{a_2}{s_2}$ iff $s_2a_1 = s_1a_2$.

If the monoid M is ordered, then $S^{-1}M$ is also ordered, by putting $\frac{a}{s} \leq \frac{a'}{s'}$ iff $as's \leq a'ss$ for some $s \in S$.

When R is a commutative semiring† and S is a submonoid of $(R, \cdot)$, we endow $S^{-1}R$ with addition given by

$$\frac{a}{s} + \frac{a'}{s'} = \frac{as' + a's}{ss'}.$$

When $(R, \cdot)$ is ordered, this is compatible with the monoid order on $S^{-1}R$, in the sense of Proposition 2.2.3.

Clearly, the localization of an idempotent (resp. proper, resp. Frobenius) semiring† is idempotent (resp. proper, resp. Frobenius).

Furthermore, if R is a domain†, and $S = R$, then we call $S^{-1}R$ the **semifield† of fractions** of R.

Remark 2.3.2. These considerations are compatible with localization and the divisible hull. For example, one can define $\vee$ on the monoid $S^{-1}M$ via

$$\frac{a}{s} \vee \frac{b}{s} = \frac{a \vee b}{s}.$$

For the divisible hull, we would use the analog of Equation 2.5.

Proposition 2.3.3. *Suppose R, R' are domains†, with a semiring† homomorphism $\varphi : R \to R'$, and suppose S is a submonoid of R. Then there is a homomorphism $\tilde{\varphi} : S^{-1}R \to \tilde{\varphi}(S)^{-1}R'$, given by $\tilde{\varphi}(\frac{a}{s}) = \frac{\varphi(a)}{\varphi(s)}$.*

Proof. This is standard. First we need to show that $\tilde{\varphi}$ is well-defined: If $\frac{a_1}{s_1} = \frac{a_2}{s_2}$ then $ss_2a_1 = ss_1a_2$ so $\varphi(s)\varphi(s_2)\varphi(a_1) = \varphi(s)\varphi(s_1)\varphi(a_2)$, implying $\varphi(s_2)\varphi(a_1) = \varphi(s_1)\varphi(a_2)$, and thus $\tilde{\varphi}(\frac{a_1}{s_1}) = \tilde{\varphi}(\frac{a_2}{s_2})$.

Now $\tilde{\varphi}\left(\frac{as_1}{ss_1}\right) = \left(\frac{a_1}{s_1}\frac{a_2}{s_2}\right) = \tilde{\varphi}\left(\frac{a_1a_2}{s_1s_2}\right) = \frac{\varphi(a_1a_2)}{\varphi(s_1s_2)} = \tilde{\varphi}\left(\frac{a_1}{s_1}\right)\tilde{\varphi}\left(\frac{a_2}{s_2}\right)$. □

2.4. **Congruences.**

A **congruence** Ω of an algebraic structure A is an equivalence relation $\equiv$ that preserves all the relevant operations and relations; we call $\equiv$ the **underlying equivalence** of Ω. One can view a congruence more formally as a subalgebra of $A \times A$. For ease of notation we write $a \equiv b$ when $(a, b) \in \Omega$. In the case of semirings†, Ω being a congruence means that $a_i \equiv b_i$ for $i = 1, 2$ implies $a_1 + a_2 \equiv b_1 + b_2$ and $a_1 a_2 \equiv b_1 b_2$.

Remark 2.4.1. We recall some key results of [21, §2]:

(i) Given a congruence Ω of an algebraic structure R, one can endow the set
$$A/\Omega := \{[a] : a \in A\}$$
of equivalence classes with the same (well-defined) algebraic structure, and the map $a \mapsto [a]$ defines an onto homomorphism $A \to A/\Omega$.

(ii) In the opposite direction, for any homomorphism $\varphi : A \to A'$, one can define a congruence Ω_φ on A by saying that
$$a \equiv_\varphi b \quad \text{iff} \quad \varphi(a) = \varphi(b).$$
We call Ω_φ the **congruence of** φ. Then φ induces a 1:1 homomorphism $\widetilde{\varphi} : A/\Omega_\varphi \to A'$, via $\widetilde{\varphi}([a]) = \varphi(a)$, for which φ factors through
$$A \to A/\Omega_\varphi \to A',$$
as indicated in [21, p. 62]. Thus the homomorphic images of A correspond to the congruences defined on A.

Definition 2.4.2. A congruence Ω on an Abelian monoid M is **cancellative** when the monoid M/Ω is cancellative. The congruence Ω on an Abelian monoid M is **power-cancellative** when the monoid M/Ω power-cancellative, i.e., if $a_1^k \equiv a_2^k$ for some $k \geq 1$ implies $a_1 \equiv a_2$.

For the moment, let $\mathcal{C}(M)$ denote a given set of cancellative and power-cancellative congruences on M. A congruence $\Omega \in \mathcal{C}(M)$ is **irreducible** (called **intersection indecomposable** in [22]) if it is not the proper intersection of two congruences in $\mathcal{C}(M)$.

Remark 2.4.3. **Prime congruences** are defined by Joo and Mincheva in [22], and are seen [22, Theorem 2.11] to be precisely the irreducible cancellative congruences. [22] goes on to classify the prime congruences of the polynomial semirings over the max-plus semifields $\mathbb{R}$ and $\mathbb{Z}$, and also over the 2-element semifield.

Lemma 2.4.4. *Any congruence Ω on a semiring† R extends to its localization $S^{-1}R$ by a cancellative submonoid S, by putting $\frac{a}{s} \equiv \frac{a'}{s'}$ when $as' \equiv a's$.*

Proof. It is standard and easy that $\equiv$ extends to the given equivalence on $S^{-1}R$, so we need check merely that the given operations are preserved: If $\frac{a_1}{s_1} \equiv \frac{a_1'}{s_1'}$ and $\frac{a_2}{s_2} \equiv \frac{a_2'}{s_2'}$, then $a_i s_i' \equiv a_i' s_i$ for $i = 1, 2$, from which it follows that $a_1 s_1' a_2 s_2' \equiv a_1' s_1 a_2' s_2$ and thus
$$\frac{a_1}{s_1}\frac{a_2}{s_2} = \frac{a_1 a_2}{s_1 s_2} \equiv \frac{a_1' a_2'}{s_1' s_2'}$$
Likewise,
$$\frac{a_1}{s_1} + \frac{a_2}{s_2} = \frac{a_1 s_2 + a_2 s_1}{s_1 s_2} \equiv \frac{a_1' s_2' + a_2' s_1'}{s_1' s_2'} = \frac{a_1'}{s_1'} + \frac{a_2'}{s_2'}.$$
□

2.5. **The function monoid and semiring†.**

The supertropical structure permits us to detect corner roots of tropical polynomials in terms of the algebraic structure, by means of ghosts. To see this most clearly, we introduce a key structure taken from universal algebra.

Definition 2.5.1. The **function monoid** $\mathrm{Fun}(S, M)$ is the set of functions from a set S to a monoid M.

The **function semiring†** $\mathrm{Fun}(\mathcal{S}, R)$ is the set of functions from a set S to a semiring† R.

When M is a monoid, $\mathrm{Fun}(S, M)$ becomes a monoid under the componentwise operation, $fg(\mathbf{a}) = f(\mathbf{a})g(\mathbf{a})$, and is cancellative when M is cancellative. If M is semi-lattice ordered, then so is $\mathrm{Fun}(S, M)$, where we define $(f \vee g)(\mathbf{a}) = f(\mathbf{a}) \vee g(\mathbf{a}) = f(\mathbf{a}) + g(\mathbf{a})$.

When R is a semiring†, $\mathrm{Fun}(\mathcal{S}, R)$ becomes a semiring† under componentwise operations, and is idempotent (resp. proper) when R is idempotent (resp. proper). When R is a semifield†, $\mathrm{Fun}(\mathcal{S}, R)$ becomes a semifield†.

Note that the semiring† $\mathrm{Fun}(\mathcal{S}, R)$ need not be bipotent even if R is bipotent, since some of the evaluations of $f+g$ might come from f and others from g. As in Proposition 2.2.3, we identify the semiring† structure of $\mathrm{Fun}(\mathcal{S}, R)$ with the semilattice operation $f \vee g = f + g$. When R is a semifield† F, we can also define

$$f \wedge g = (f^{-1} + g^{-1})^{-1}. \tag{2.7}$$

Remark 2.5.2. Viewing F as an ordered group, one sees easily that $(f \wedge g)(\mathbf{a}) = \min\{f(\mathbf{a}), g(\mathbf{a})\}$, and this is how we think of it, although (2.7) is easier to handle formally.

Often $\mathcal{S}$ is taken to be $F^{(n)}$, the Cartesian product of n copies of F. But using $\mathcal{S} \subseteq F^{(n)}$ enables one to lay the foundations of algebraic geometry.

For any subset $\mathcal{S}' \subseteq \mathcal{S}$, there is a natural homomorphism

$$\mathrm{Fun}(S, F) \to \mathrm{Fun}(\mathcal{S}', F),$$

given by restriction of functions.

2.6. **Archimedean monoids.**

Definition 2.6.1. A partially ordered monoid $(G, \cdot)$ is called **archimedean** if $a^{\mathbb{Z}} \leq b$ implies that $a = 1$.

A semiring† $(R, +, \cdot)$ is **archimedean** if it is archimedean as a partially ordered monoid.

This reduces to the usual definition when G is totally ordered, but we need this more general version, because of the next example.

Example 2.6.2. *The semifield*† $\mathbb{N}$ *is ordered and archimedean, as are* $(\mathbb{Q}, +)$ *and* $(\mathbb{R}, +)$.

Remark 2.6.3. If R is archimedean then $\mathrm{Fun}(\mathcal{S}, R)$ also is archimedean.

3. The tropical environment

We would like to fit this all into tropical mathematics, mainly for affine varieties defined by a finite number of polynomials. Tropical mathematics involves using the Puiseux valuation to pass from the field of Puiseux series to its value group $(\mathbb{Q}, +)$ (or $(\mathbb{R}, +)$), which in turn is viewed as the max-plus algebra. But this transition does not reflect the valuation theory for Puiseux series having lowest order terms of the same degree; their sums do not necessarily have lowest order terms of this degree. Thus we should consider alternative structures in which $a \vee a \neq a$ – i.e., monoids which are ordered by sets which are not quite semi-lattices.

3.1. **The standard supertropical semifield†.**

We are ready to bring in the algebraic structure that reflects the properties of the Puiseux valuation.

Definition 3.1.1. A **ν-domain**† is a quadruple $(R, \mathcal{T}, \nu, \mathcal{G})$ where R is a semiring† and $\mathcal{T} \subset R$ is a cancellative multiplicative submonoid and $\mathcal{G} \lhd R$ is endowed with a partial order, together with an idempotent homomorphism $\nu : R \to \mathcal{G}$, with $\nu|_{\mathcal{T}}$ onto, satisfying the conditions:

$$a + b = a \quad \text{whenever} \quad \nu(a) > \nu(b).$$

$$a + b = \nu(a) \quad \text{whenever} \quad \nu(a) = \nu(b).$$

$\mathcal{T}$ is called the **tangible submonoid** of R. $\mathcal{G}$ is called the **ghost ideal**.

(In practice, ν is induced from an m-valuation of a field, for example the Puiseux valuation.) The condition $\nu|_{\mathcal{T}}$ onto is introduced to enable us to move back and forth between tangible and ghost elements.

Remark 3.1.2. This definition is quite close to that of [13], where $\mathcal{T}$ would be the tangible elements and $\mathcal{G}$ the ghost ideal $(1+1)R$. The main differences are as follows:

- Here the target $\mathcal{G}$ need not be ordered, so that we can handle semirings of polynomials directly.
- We do not require a zero element.

We write a^ν for $\nu(a)$, for $a \in R$. We write $a \cong_\nu b$ if $a^\nu = b^\nu$, and say that a and b are **ν-equivalent**. Likewise we write $a \geq_\nu b$ (resp. $a >_\nu b$) if $a^\nu \geq b^\nu$ (resp. $a^\nu > b^\nu$).

Remark 3.1.3. Any semiring† homomorphism $\varphi : R \to R' = (R', \mathcal{T}', \mathcal{G}', \nu')$ of ν-domains† satisfies

$$\varphi(a^\nu) = \varphi(a + a) = \varphi(a) + \varphi(a) = \varphi(a)^{\nu'}$$

for all a in R.

Strictly speaking, the supertropical ν-domains† can never be semifields† since the ghost elements are not invertible. Accordingly, we frame the next definition:

Definition 3.1.4. A ν-domain† R is **ν-bipotent** if $a+b \in \{a, b, a^\nu\}$ for all $a, b \in R$. A **ν-semifield†** $\mathbb{S}$ is a ν-domain† for which the tangible submonoid $\mathcal{T}$ is an Abelian group.

Example 3.1.5. *In the case where $F = \mathcal{T} = \mathcal{G}$, with $\nu = 1_F$, we are back to the semifield† theory. We call this the **degenerate case** Although we are not interested in this case, it shows us how to generalize to the ν-theory. Namely, we must "lift" the usual theory from $\mathcal{G}$ to F.*

Example 3.1.6. *Any supertropical semifield† in the sense of* [17] *is a ν-semifield†, in the sense of this paper.*

The most important example, which we denote as $\mathscr{R}$, is defined by taking the tangible submonoid to be $(\mathbb{R}, +)$, together with another ghost copy.

Here is our main example, but we will need our more general definition when we consider functions, and in particular, polynomial functions.

Example 3.1.7. *Given a group M and an ordered group $\mathcal{G}$ with a group isomorphism $\nu : M \to \mathcal{G}$, we write a^ν for $\nu(a)$. The **standard supertropical** monoid R is the disjoint union $\mathcal{T} \cup \mathcal{G}$ where $\mathcal{T}$ is taken to be M, made into a group by starting with the given products on M and $\mathcal{G}$, and defining ab^ν and $a^\nu b$ to be $(ab)^\nu$ for $a, b \in M$.*

*We extend v to the **ghost map** $\nu : R \to \mathcal{G}$ by taking $\nu|_M = v$ and $\nu_{\mathcal{G}}$ to be the identity on $\mathcal{G}$.*

*We make R into a semiring†, called the **standard supertropical semifield†**, by defining*

$$a+b = \begin{cases} a \text{ for } a >_\nu b; \\ b \text{ for } a <_\nu b; \\ a^\nu \text{ for } a \cong_\nu b. \end{cases}$$

In order for Condition (iv) of Definition 2.2.1 to hold, we need the valuation ν to be nontrivial, in the sense that there exists $a \in R$ such that $a <_\nu \mathbf{1}$.

R is ν-bipotent. R is never additively cancellative, since

$$a + a^\nu = a^\nu = a^\nu + a^\nu.$$

The standard supertropical semifield† R is a cover of the max-plus algebra of $\mathcal{G}$, in which we can "resolve" additive idempotents, in the sense that $a + a = a^\nu$ instead of $a + a = a$.

Although we are interested in $M = \mathcal{T}$, the ordered group $\mathcal{G}$ itself can be viewed as an idempotent semiring† with unit element $\mathbf{1}_R^\nu$ (instead of $\mathbf{1}_R$). Thus, we obtain ν-versions of the previous notions by lifting from $\mathcal{G}$.

The ghost ideal $\mathcal{G}$ is to be treated much the same way that one treats the 0 element in commutative algebra. The standard supertropical semifield† works well with linear algebra, as seen for example in [18], providing many of the analogs to the classical Cayley-Hamilton-Frobenius theory, but our interest here will be in its geometric significance.

3.1.1. *ν-Localization.*

If $R = (R, \mathcal{T}, \mathcal{G}, \nu)$ is a ν-domain†, then we call $F := \mathcal{T}^{-1}R$ the **ν-semifield† of fractions** $\mathrm{Frac}_\nu R$ of R.

Lemma 3.1.8. *$\mathrm{Frac}_\nu R$ is a ν-semifield† in the obvious way, where $\mathcal{T}^{-1}\mathcal{T}$ is the group of tangible elements.*

Proof. Define $\nu(\frac{r}{s}) = \frac{r^\nu}{s}$. □

3.2. **The (*)-operation.**

Strictly speaking, since a ν-semifield† is not a semifield†, we need to modify our definition of inverse. Considering $\mathcal{G}$ as a group in its own right, with unit element 1^ν, we define a^* to be the inverse of a in the appropriate group to which it belongs ($\mathcal{T}$ or $\mathcal{G}$). Specifically, we have:

Remark 3.2.1. Any supertropical semifield† $(F, \mathcal{T}, \mathcal{G}, \nu)$ has a multiplicative monoid automorphism $(*)$ of order 2 given by

$$a^* = a^{-1}, \qquad (a^\nu)^* = (a^{-1})^\nu, \qquad a \in \mathcal{T}.$$

But $(*)$ reverses the partial order induced by ν, and thus does not preserve addition.

Remark 3.2.2. There is ambiguity in defining $a \wedge b$ in a supertropical semifield†, since we can only compare ν-values. So we define

$$a \wedge b = (a^* + b^*)^*. \tag{3.1}$$

But now $a \wedge a = a^\nu$.

3.2.1. *The ν-norm.*

Since the element $\mathbf{1}_F$ plays an important role, the following notion will be useful.

Definition 3.2.3. Suppose $F = (F, \mathcal{T}, \mathcal{G}, \nu)$ is a supertropical semifield†. The **ν-norm** $|a|$ of an element $a \in F$ is $a + a^*$.

In the degenerate case of Example 3.1.5, $|a| = a + a^{-1}$. In the max-plus semifield† this is the parallel of the usual absolute value of a.

Remark 3.2.4. (i) $|a|^2 = \mathbf{1}^\nu + |a^2| \geq_\nu \mathbf{1}$.

(ii) $|a| \cong_\nu \mathbf{1}$ iff $a \cong_\nu \mathbf{1}$.

3.3. **The function ν-domain†.**

This supertropical structure also permits us to detect corner roots of tropical polynomials in terms of the algebraic structure by means of ghosts, once we provide a ν-structure for the function semiring† of Definition 2.5.1.

Example 3.3.1. *If $(R, \mathcal{T}, \mathcal{G}, \nu)$ is a ν-domain†, then* $\mathrm{Fun}(\mathcal{S}, R)$ *is a ν-domain† has the induced map* $\tilde{\nu} : \mathrm{Fun}(\mathcal{S}, R) \to \mathrm{Fun}(\mathcal{S}, \mathcal{G})$ *given by*

$$\tilde{\nu}(f)(\mathbf{a}) = \nu(f(\mathbf{a})).$$

Here the ghosts are $\mathrm{Fun}(\mathcal{S}, \mathcal{G})$.

We also extend the $(*)$-map to $\mathrm{Fun}(\mathcal{S}, F)$*:*

For F a supertropical semifield† and $f \in \mathrm{Fun}(\mathcal{S}, F)$, define f^* via $f^*(a) = f(a)^*$ for all $a \in F$.

By Remark 3.2.1, the map $f \mapsto f^*$ defines a monoid automorphism of $\mathrm{Fun}(\mathcal{S}, F)$ of order 2. In this way, we have the following lattice operations on $\mathrm{Fun}(\mathcal{S}, F)$:

$$f \vee g = f + g; \qquad f \wedge g = (f^* + g^*)^*. \tag{3.2}$$

Definition 3.3.2. For any semitropical semifield† $F = (F, \mathcal{T}, \mathcal{G}, \nu)$ and any subset $A \subseteq \mathrm{Fun}(S, F)$, A^+ denotes $\{f \in A : f >_\nu 1\}$.

Remark 3.3.3. If R is a ν-domain†, then so is R^+.

3.3.1. *Polynomials and Laurent polynomials.*

$\Lambda = \{\lambda_1, \dots, \lambda_n\}$ always denotes a finite commuting set of indeterminates; often $n = 1$, and we have a single indeterminate λ.

Given a semiring† R, we have the polynomial semiring† $R[\Lambda]$. Just as in [17], we view polynomials in $R[\Lambda]$ as functions. More precisely, for any subset $\mathcal{S} \subseteq R^{(n)}$, there is a natural semiring† homomorphism

$$\psi : R[\Lambda] \to \mathrm{Fun}(\mathcal{S}, R), \tag{3.3}$$

obtained by viewing a polynomial as a function on $\mathcal{S}$, and we write $.\overline{R[\Lambda]}$ for $\psi(R[\Lambda]) \subset \mathrm{Fun}(\mathcal{S}, R)$. (Likewise for $\mathcal{T}[\Lambda]$ and $\mathcal{G}[\Lambda]$.

To ease notation, we still write a typical element of $\overline{R[\Lambda]}$ as $\left(\sum a_i \lambda_1^{i_1} \dots \lambda_n^{i_n}\right)$. Any m-valuation $\nu : R \to \mathcal{G}$ extends to a map $\tilde{\nu} : \overline{R[\Lambda]} \to \overline{\mathcal{G}[\Lambda]}$ given by

$$\tilde{\nu}\left(\sum a_i \lambda_1^{i_1} \dots \lambda_n^{i_n}\right) = \sum \nu(a_i) \lambda_1^{i_1} \dots \lambda_n^{i_n}.$$

Remark 3.3.4. If $R = (R, \mathcal{T}, \mathcal{G}, \nu)$ is a ν-domain†, then so is $(\overline{R[\Lambda]}, \overline{\mathcal{T}[\Lambda]}, \overline{\mathcal{G}[\Lambda]}, \nu)$ (where $\lambda^\nu = \mathbf{1}_R^\nu \lambda$).

This definition is concise, but we should note the difficulty that two polynomials, one tangible, and one non-tangible, can describe the same function, for example $\lambda^2 + 6$ and $(\lambda + 3)^2 = \lambda^2 + 3\lambda^\nu + 6$. One can remedy this formally by identifying polynomials agreeing on a dense subset of a variety.

When R is a ν-semifield†, the same analysis is applicable to the ν-domain† of Laurent polynomials $\overline{R[\Lambda, \Lambda^{-1}]}$, since the homomorphism $\lambda_i \mapsto a_i$ then sends $\lambda_i^{-1} \mapsto a_i^{-1}$. But the algebraic structure of choice in this paper is the following (and its ν-analog described below):

Definition 3.3.5. When $(F, \mathcal{T}, \mathcal{G}, \nu)$ is a supertropical ν-semifield†, $\overline{F(\Lambda)}$ denotes $\mathrm{Frac}_\nu \overline{F[\Lambda]}$, the ν-semifield† of fractions of $\overline{F[\Lambda]}$ obtained by localizing at $\overline{\mathcal{T}[\Lambda]}$. We call $\overline{F(\Lambda)}$ the **ν-semifield† of rational functions** (over F); this is our main subject of investigation.

Viewed in $\mathrm{Fun}(F^{(n)}, F)$, $\overline{F(\Lambda)}$ is contained in $\{fg^* : f, g \in \overline{F[\Lambda]}\}$, and they behave similarly. We use the former ($\overline{F(\Lambda)}$) since it is intuitively clearer, although probably the latter may be technically superior.

(Likewise, we could also define $\overline{F[\Lambda]}_{\mathrm{rat}}$ where the exponents of the indeterminates λ_i are taken to be arbitrary rational numbers; we could define substitution homomorphisms when $(F, \mathcal{T}, \mathcal{G}, \nu)$ is power-cancellative and divisible.)

These can all be viewed as elementary sentences in the appropriate language, so model theory is applicable to polynomials and their (tropical) roots, to be considered shortly.

Remark 3.3.6. Here are the semilattice operations on $\overline{F(\Lambda)}$:

Suppose $f = \frac{h}{g}$ and $f' = \frac{h'}{g'}$. Then

$$f \vee f' = f + f' = \frac{hg' + gh'}{gg'}; \qquad f \wedge f' = ff'(f + f')^* = \frac{hh'}{gg'} gg'(hg' + gh')^* = hh'(hg' + gh')^*. \tag{3.4}$$

When $hg' + gh'$ is tangible, we have

$$f \wedge f' = \frac{hh'}{hg' + gh'}. \tag{3.5}$$

3.4. **Frobenius and archimedean properties.**

Definition 3.4.1. A ν-domain† R is **ν-Frobenius** if it satisfies the identity

$$(a + b)^n \cong_\nu a^n + b^n, \qquad \forall a, b \in R.$$

R is **ν-archimedean** if if $a^{\mathbb{Z}} \leq_\nu b$ implies that $a \cong_\nu 1$.

Proposition 3.4.2. *If R is a ν-Frobenius ν-domain†, then so is* $\mathrm{Fun}(\mathcal{S}, R)$.

Proof. Pointwise verification. □

Proposition 3.4.3. *If the idempotent semifield† F is ν-archimedean, then $Fun(F^{(n)}, F)$ is ν-archimedean.*

Proof. Let $f, g \in Fun(F^{(n)}, \mathcal{G})$ such that $f^{\mathbb{Z}} \leq_\nu g$. If $\mathbf{a} \in F^{(n)}$ then by our assumption $f(\mathbf{a})^k \leq_\nu g(\mathbf{a})$ for all $k \in \mathbb{Z}$. Since F is ν-archimedean and $f(\mathbf{a}), g(\mathbf{a}) \in F$, we have $f(\mathbf{a}) = 1$. Since this holds for any $\mathbf{a} \in F^{(n)}$ we have $f(F^{(n)}) \cong_\nu 1$, implying $f \cong_\nu 1$. □

Lemma 3.4.4. *If a supertropical semifield F is ν-Frobenius, then k-th roots in $Fun(F^{(n)}, F)$ are unique.*

Proof. It is enough to check this pointwise in F. Suppose $a^k = b^k$. Then $(a + b)^k \cong_\nu a^k + b^k \in \mathcal{G}$, implying $a + b \in \mathcal{G}$. If moreover, a^k is tangible then so is a and b, so $a = b$. But if $a^k \in \mathcal{G}$, then a and b are both ghost, so $a = b$. □

Remark 3.4.5. Let $(\mathcal{S}, \cdot, +)$ be a divisible semifield†. By Lemma 3.4.4, we can uniquely define any rational power of the elements of $\mathcal{S}$. In this way, $\mathcal{S}$ becomes a vector space over $\mathbb{Q}$, rewriting the multiplicative operation $\cdot$ on $\mathcal{S}$ as addition and defining

$$(m/n) \cdot \alpha = \alpha^{\frac{m}{n}}. \tag{3.6}$$

Furthermore, when $\mathcal{S}$ is complete, we can define real powers as limits of rational powers, and $\mathcal{S}$ becomes a vector space over $\mathbb{R}$. In this way we can apply linear algebra techniques to $(\mathcal{S}, \cdot)$.

4. Kernels and ν-kernels

Although the theory of congruences applies generally in universal algebra and in particular for a semiring† R, congruences are difficult to work with since they involve subsets of $R \times R$ rather than R itself. When R is a semifield† $\mathcal{S}$, one can get around this by switching the roles of addition and multiplication. Although we are interested in the ν-structure, we present the definitions without ν, in order to get to the underlying ideas more quickly.

4.1. ν-Congruences and ν-kernels of ν-semifields†.

Definition 4.1.1. A **kernel** of a semiring† $\mathcal{S}$ is a subgroup K which is **convex** in the sense that if $a, b \in K$ and $\alpha, \beta \in \mathcal{S}$ with $\alpha + \beta = \mathbf{1}_{\mathcal{S}}$, then $\alpha a + \beta b \in K$.

Note 4.1.2. Usually one takes $\mathcal{S}$ to be a semifield†, which is our assumption through the remainder of this subsection, but we frame the definition a bit more generally in order to be able later to handle ν-semifields†, which technically are not semifields†. In some of the literature semifields† are not required to be commutative, and then K is required to be a normal subgroup; we forego this generality.

Following [37, Proposition (1.1)], we have the following key correspondence.

Proposition 4.1.3. *If Ω is a congruence on a semifield† $\mathcal{S}$, then $K_\Omega = \{a \in \mathcal{S} : a \equiv 1\}$ is a kernel. Conversely, any kernel K of $\mathcal{S}$ defines a congruence ρ_K according to* [12, Definition 3.1]*, i.e., $a \equiv b$ iff $\frac{a}{b} \equiv 1$. If $\mathcal{S}$ is the semifield† of the lattice-ordered group G, then the semifield† $\mathcal{S}/\rho_K$ is the semifield† of the lattice-ordered group G/K.*

A quick proof that Ω is indeed a congruence is given in [38, Proposition 1.1].

We need to make a slight modification in order to take the ν-structure into account. Throughout, $R = (R, \mathcal{T}, \mathcal{G}, \nu)$ is a ν-domain†. Any congruence of $\mathcal{G}$ also can be viewed as a congruence of R, but we also want to bring $\mathcal{T}$ into play. The next observation is due to Izhakian.

Remark 4.1.4. Any congruence Ω satisfies the property that if $(a, b) \in \Omega$ then $(a^\nu, b^\nu) = (1^\nu, 1^\nu)(a, b) \in \Omega$.

Definition 4.1.5. A **ν-congruence** is a congruence Ω for which $(a, b) \in \Omega$ iff $(a^\nu, b^\nu) \in \Omega$. We write $a_1 \equiv_\nu a_2$ when $a_1^\nu \equiv a_2^\nu$.

A **ν-kernel** of a ν-semifield† F is a subgroup K which is **ν-convex** in the sense that if $a, b \in K$ and $\alpha, \beta \in F$ with $\alpha + \beta \cong_\nu \mathbf{1}_F$, then $\alpha a + \beta b \in K$.

Remark 4.1.6. There is a natural correspondence between ν-congruences (resp. ν-kernels) of R and congruences (resp. kernels) of $\mathcal{G}$, given as follows:

Any ν-congruence $\Omega = \{(a, b) : a, b \in R\}$ of R defines a congruence $\Omega^\nu = \{(a^\nu, b^\nu) : (a, b) \in \Omega\}$ of $\mathcal{G}$. Conversely, if Ω_ν is a congruence of $\mathcal{G}$, then $\nu^{-1}(\Omega_\nu)$ is a ν-congruence of R.

Any ν-kernel $\mathcal{K}$ of a ν-semifield† F defines a kernel $\mathcal{K}^\nu$ of $\mathcal{G}$. Conversely, if $\mathcal{K}$ is a kernel of $\mathcal{G}$, then $\nu^{-1}(\mathcal{K})$ is a ν-kernel of F.

If $\mathcal{K} = \mathcal{K}_\Omega$, then $\nu^{-1}(\mathcal{K}) = \mathcal{K}_{\nu^{-1}(\Omega)}$.

Thus, we will quote results from the literature about kernels and use their analogs for ν-kernels.

Theorem 4.1.7. *Given a ν-congruence Ω on a ν-semifield† F, corresponding to the ν-kernel K of F, the factor semiring† $\bar{F} = F/\Omega$ is a ν-semifield† with respect to the idempotent map $\bar{\nu} : \bar{F} \to \overline{\mathcal{G}} = \mathcal{G}/\nu(K)$ defined by $\bar{\nu}(Ka) = Ka^\nu$.*

Proof. If $a \cong_\nu b$ then $ab^* \in K$, implying $Ka = Kb$. This implies $\bar{\nu}$ is well-defined, and its target is the semifield† $\overline{\mathcal{G}}$. □

Writing $\bar{a}$ for the image of a in $\bar{R}$, we write $a \equiv_\nu b$ to denote that $\bar{a} \cong_\nu \bar{b}$, or equivalently $\overline{a}\overline{b}^* \cong_\nu \bar{\mathbf{1}}_F$.

Remark 4.1.8.
(i) Given a ν-congruence Ω on a ν-semifield† F, we define $K_\Omega = \{a \in F : a \equiv_\nu 1\}$. Conversely, given a ν-kernel K of F, we define the ν-congruence Ω on F by $a \equiv b$ iff $ab^* \equiv_\nu \mathbf{1}_F$.
(ii) [38, Corollary 1.1], [37, Property 2.4] Any ν-kernel K is ν-convex, in the sense that if $a \leq_\nu b \leq_\nu c$ with $a, c \in K$, then $b \in K$. This is seen by passing to the factor semiring† F/Ω (Ω as in (i)) and applying Lemma 2.1.6.
(iii) [38, Proposition 2.3]. If $|a| \in K$, a ν-kernel, then $a \in K$.
(iv) [38] The product $K_1K_2 = \{ab \ : \ a \in K_1, b \in K_2\}$ of two ν-kernels is a ν-kernel, in fact the smallest ν-kernel containing $K_1 \cup K_2$. (This follows at once from (ii).)
(v) The intersection of ν-kernels is a ν-kernel. Thus, for any set $S \subset F$ we can define the ν-kernel $\langle S \rangle$ **generated by** S to be the intersection of all ν-kernels containing S.
(vi) [38, Theorem 3.5]. Any ν-kernel generated by a finite set $\{s_1, \dots, s_m\}$ is in fact generated by the single element $\sum_{i=1}^m (|s_i|)$. (Follows easily from (iii).)
(vii) The ν-kernel generated by $a \in F$ is just the set of finite sums $\{\sum_i b_i a^i : \ b_i \in F, \ \sum b_i \cong_\nu 1\}$.
(viii) [12, Theorem 3.8]. If K is a ν-kernel of a ν-semifield† F and the ν-semifield† F/K is ν-idempotent, then K is a sub-ν-semifield† of F. (This is because for $a, b \in K$ the image of $a + b$ is ν-equivalent to $1K + 1K = 1K$.)

(ix) Let K be a ν-kernel of a ν-semifield† F. For every $a \in \mathcal{T}$, if $a^n \in K$ for some $n \in \mathbb{N}$ then $a \in K$. (Indeed, passing to F/K, it suffices to note by the analog of Lemma 2.2.4 that if $a^n = 1$ then $a = 1$.)
(x) Any ν-kernel of a ν-kernel is a ν-kernel.

Proposition 4.1.9. [36, Theorem (2.2.5)] *The lattice of ν-kernels of an idempotent ν-semifield† is a complete distributive lattice and satisfies the infinite distributive law.*

Theorem 4.1.10. [12, Theorems (3.4) and (3.5)] *Let $\phi : F_1 \to F_2$ be a ν-semifield† homomorphism. Then the following hold:*

(1) *For any ν-kernel L of F_1, $\phi(L)$ is a ν-kernel of $\phi(F_1)$.*
(2) *For a ν-kernel K of $\phi(F_1)$, $\phi^{-1}(K)$ is a ν-kernel of F_1.*

In particular, $\phi^{-1}(1)$ is a ν-kernel.

Corollary 4.1.11. *There is an injection $\mathbb{S}/(K_1 \cap K_2 \cap \cdots \cap K_t) \hookrightarrow \prod_{i=1}^t \mathbb{S}/K_i$, for any ν-kernels K_i of a ν-semifield† $\mathbb{S}$, induced by the map $f \mapsto (fK_i)$.*

We have the fundamental isomorphism theorems.

Theorem 4.1.12. [40] *Let $\mathcal{S}$ be an idempotent ν-semifield† and K, L be ν-kernels of $\mathcal{S}$.*

(1) *If $\mathcal{U}$ is a sub-ν-semifield† of $\mathcal{S}$, then $\mathcal{U} \cap K$ is a ν-kernel of $\mathcal{U}$, and K a ν-kernel of the sub-ν-semifield† $\mathcal{U}K = \{u \cdot k \ : \ u \in \mathcal{U}, \ k \in K\}$ of $\mathcal{S}$, and one has the isomorphism*
$$\mathcal{U}/(\mathcal{U} \cap K) \cong \mathcal{U}K/K.$$
(2) *$L \cap K$ is a ν-kernel of L and K a ν-kernel of LK, and the group isomorphism*
$$L/(L \cap K) \cong LK/K$$
is a ν-semifield† isomorphism.
(3) *If $L \subseteq K$, then K/L is a ν-kernel of $\mathcal{S}/L$ and one has the ν-semifield† isomorphism*
$$\mathcal{S}/K \cong (\mathcal{S}/L)/(K/L).$$

Theorem 4.1.13. *Let L be a ν-kernel of a ν-semifield† F. Every ν-kernel of F/L has the form K/L for some uniquely determined ν-kernel $K \supseteq L$, yielding a lattice isomorphism*
$$\{\text{Kernels of } F/L\} \to \{\text{Kernels of } F \text{ containing } L\}$$
given by $K/L \mapsto K$.

Proof. From the theory of groups, there is such a bijection for subgroups. We note by Theorem 4.1.10 that the homomorphic image and pre-image of a ν-kernel are ν-kernels. □

Definition 4.1.14. A ν-semifield† which contains no ν-kernels but the trivial ones, $\{1, 1^\nu\}$ and itself, is called **simple**.

Remark 4.1.15. A ν-kernel K of a ν-semifield† F is a maximal ν-kernel if and only if the ν-semifield† F/K is simple.

4.1.1. *Principal ν-kernels.*

Definition 4.1.16. A ν-kernel K is said to be **finitely generated** if $K = \langle S \rangle$ where S is a finite set. If $K = \langle a \rangle$ for some $a \in F$, then K is called a **principal ν-kernel**.

$\mathcal{P}(K)$ denotes the set of principal ν-subkernels of a ν-kernel K. $\mathcal{P}(K)$ turns out to be a lattice, and one of the keys of our study.

Lemma 4.1.17. [37, Property 2.3] *Let K be a ν-kernel of an idempotent ν-semifield† $\mathcal{S}$. Then for $a, b \in \mathcal{S}$,*
$$|a| \in K \quad \text{or} \quad |a| + b \in K \quad \Rightarrow a \in K. \tag{4.1}$$

Proposition 4.1.18. [37, Proposition (3.1)]
$$\langle a \rangle = \{x \in \mathcal{S} \ : \ \exists n \in \mathbb{N} \quad \text{such that} \quad a^{-n} \leq x \leq a^n\}. \tag{4.2}$$

Corollary 4.1.19. [37] *$\langle a \rangle = \langle a^k \rangle$, for any element a of $\mathcal{S}$ and any $0 \neq k \in \mathbb{Z}$.*

Corollary 4.1.20. [37] *$\langle a \rangle = \langle |a| \rangle$, for any element a of $\mathcal{S}$.*

A direct consequence of Proposition 4.1.18 and Corollary 4.1.20 is:

Corollary 4.1.21. *For any $a \in \mathcal{S}$,*

$$\langle a \rangle = \{x \in \mathcal{S} \ : \ \exists n \in \mathbb{N} \text{ such that } |a|^{-n} \leq_\nu x \leq_\nu |a|^n\}.$$

Definition 4.1.22. A ν-semifield† is said to be **finitely generated** if it is finitely generated as a ν-kernel. If $\mathcal{S} = \langle a \rangle$ for some $a \in \mathcal{S}$, then a is called a **generator** of $\mathcal{S}$ and $\mathcal{S}$ is said to be a **principal ν-semifield†**.

Corollary 4.1.23. *Every nontrivial principal ν-semifield† $\mathcal{S}$ has a generator $a > 1$.*

Proof. Let $u \in \mathcal{S} \setminus \{1\}$ be a generator of $\mathcal{S}$. By Lemma 4.1.17, the element $|u|$ is also a generator of $\mathcal{S}$ which yields that the element $a = |u|^2 = u^2 + u^{-2} + 1 \geq 1$ is a generator of $\mathcal{S}$ too, by Proposition 4.1.18. But $a \neq 1$, so $a > 1$. □

Theorem 4.1.24. *If an idempotent ν-semifield† $\mathcal{S}$ has a finite number of generators $a_1, \ldots, a_n$, then $\mathcal{S}$ is a principal ν-semifield†, generated by $a = |a_1| + \cdots + |a_n|$.*

Proof. By Lemma 4.1.17, $a_1, ..., a_n$ are contained in the ν-kernel $\langle a \rangle \subseteq \mathcal{S}$; hence, $\mathcal{S} = \langle a \rangle$ as desired. □

Proposition 4.1.25. *Every idempotent ν-archimedean ν-semifield† $\mathcal{S}$ is simple.*

Proof. We may assume that $\mathcal{S} \neq \{1\}$. Take $a \in \mathcal{S}$ such that $a > 1$. Since $\mathcal{S}$ is ν-archimedean, for every $b \in \mathcal{S}$ there exists $m \in \mathbb{N}$ such that $a^{-m} \leq_\nu b \leq_\nu a^m$, so $b \in \langle a \rangle$ by Proposition 4.1.18. Thus $\langle a \rangle = \mathcal{S}$ and our claim is proved. □

Notation 4.1.26. By Proposition 4.1.25, $\mathcal{S} = \langle \alpha \rangle$ for each $\alpha \neq \{1\}$.

The ν-kernel $\langle \mathcal{S} \rangle$ in $\mathcal{S}(\Lambda)$ is somewhat bigger, containing all functions $\alpha f + \beta g$ where $f + g$ is the constant function 1. This has a significant role, discussed in §8.1.1 below, where we describe $\mathcal{S}(\Lambda)$ more explicitly.

Proposition 4.1.27. *For any $\gamma_1, ..., \gamma_n \in \mathcal{S}$ the ν-kernel $\langle \frac{\lambda_1}{\gamma_1}, ..., \frac{\lambda_n}{\gamma_n} \rangle$ is a maximal ν-kernel of $\mathcal{S}(\Lambda)$.*

Proof. The quotient is isomorphic to $\mathcal{S}$, which is simple. □

Here are more properties of ν-kernels of an idempotent ν-semifield† and their generators.

Proposition 4.1.28. [36, Theorem 2.2.4(d)] *For any $V, W \subset \mathcal{S}$, any ν-kernel K of $\mathcal{S}$ and $f, g \in \mathcal{S}$ the following statements hold:*

(i) $\langle V \rangle \langle W \rangle = \langle V \cup W \rangle = \langle \{|x| + |y| : x \in V, y \in W\} \rangle$.
(ii) $\langle V \rangle \cap \langle W \rangle = \langle \{|x| \wedge |y| : x \in V, \ y \in W\} \rangle$.
(iii) $\langle K, f \rangle \cap \langle K, g \rangle = K \langle |f| \wedge |g| \rangle$.

Lemma 4.1.29.

$$|g||h| \in K \quad \Leftrightarrow \quad |g| + |h| \in K. \tag{4.3}$$

for any ν-kernel K of an idempotent ν-semifield†.

Proof. $|g| \leq |g||h|$ and $|h| \leq |g||h|$ since $|g|, |h| \geq 1$, and thus

$$|g| + |h| = \sup(|g|, |h|) \leq |g||h|.$$

On the other hand,

$$(|g| + |h|)^2 = |g|^2 + |g||h| + |h|^2 \geq |g||h|,$$

so, by Corollary 4.1.21, $\langle |g||h| \rangle = \langle |g| + |h| \rangle$ and (4.3) follows. □

Thus, we have the following fact, essentially garnered from [36].

Theorem 4.1.30. *Let $\mathcal{S}$ be an idempotent ν-semifield†. Then the intersection and product of two principal ν-kernels are principal ν-kernels. Namely, for every $f, g \in \mathcal{S}$*

$$\langle f \rangle \cap \langle g \rangle = \langle |f| \wedge |g| \rangle \ ; \ \langle f \rangle \langle g \rangle = \langle |f||g| \rangle. \tag{4.4}$$

Proof. Taking $V = \{f\}$ and $W = \{g\}$ in Proposition 4.1.28 yields the first equality, whereas for the second equality, Lemma 4.1.29 yields

$$\langle f \rangle \langle g \rangle = \langle |f| + |g| \rangle,$$

and again we apply the proposition. □

Corollary 4.1.31. *The set of principal ν-kernels of an idempotent ν-semifield† forms a sublattice of the lattice of ν-kernels (i.e., a lattice with respect to intersection and multiplication).*

Corollary 4.1.32. *For any generator a of a ν-semifield† F, $\overline{F(\Lambda)} = \langle a\rangle \prod_{i=1}^n \langle \lambda_i\rangle$, and $\overline{F(\Lambda)}$ is a principal ν-semifield† with generator $\sum_{i=1}^n |\lambda_i| + |a|$.*

Proof. By Corollary 4.1.31 it is enough to prove that any monomial f in $\overline{F[\Lambda]}$ belongs to $\langle \sum_{i=1}^n |\lambda_i| + |a|\rangle$, which follows from Theorem 4.1.24. □

4.2. **Digression: Fractional ν-kernels.**

Having a way of transferring information to and from semifields†, let us introduce fractional ν-kernels, to be able to pass to ν-domains†. This should be a useful tool in the further study of ν-kernels in tropical mathematics.

Definition 4.2.1. A **fractional ν-kernel** K on a ν-domain† R is a ν-kernel on the ν-semifield† $\mathrm{Frac}_\nu R$.

Remark 4.2.2. Extending Remark 4.1.8, given a ν-congruence Ω on a ν-domain† R with ν-semifield† of fractions F, we can define the fractional ν-kernel $K_\Omega = \{ab^* \in F : a, b \in R \text{ and } a \equiv_\nu b\}$. (Or, in other words, extending Ω to F as in Lemma 2.4.4, we require that $ab^* \equiv_\nu \mathbf{1}_F$.)

A ν-congruence Ω on R is called **ν-cancellative** if $a_1 b \equiv_\nu a_2 b$ for some $b \in R$ implies $a_1 \equiv_\nu a_2$.

Theorem 4.2.3. *There is a 1:1 correspondence between ν-cancellative congruences of a ν-domain† R and fractional ν-kernels, given by $\Omega \to K_\Omega$. Any homomorphism of ν-domains† $R \to R'$ (for R' an arbitrary ν-domain†) gives rise to a fractional ν-kernel, and Ω is the congruence corresponding to $R \mapsto R/K$.*

Proof. We extend Ω to F as in Lemma 2.4.4. As noted in [12, Theorem 3.2], K_Ω is clearly a congruence, since $\alpha a + \beta b \equiv \alpha + \beta \cong_\nu \mathbf{1}_F$ for all $a, b \in K_\Omega$.

In the other direction, we restrict this to R.

Given a homomorphism $R \to R'$, we compose it with the injection of R' into its ν-semifield† of fractions F', and then extend this naturally to a homomorphism of semifields† $F \to F'$, thereby obtaining a ν-kernel. □

5. The supertropical connection: Corner loci and K-varieties

In this section we apply our theory to tropical geometry. Let us recall the basic notions concerning supertropical varieties. Throughout, F denotes a ν-semifield†.

5.1. **Corner loci.**

In [17, Section (5.2)] Izhakian and Rowen have generalized the notion of (tangible) corner root to $F[\Lambda]$ (over a supertropical ν-semifield† F) as follows:

Definition 5.1.1. Suppose $f \in F[\Lambda]$ is a supertropical polynomial, written $f = \sum_{i=1}^k f_i$ where each f_i is a monomial. A point $\mathbf{a} \in F^{(n)}$ is said to be a **ghost root** of f if $f(\mathbf{a}) \in \mathcal{G}$, i.e., if f has a ghost value at $\mathbf{a}$.

This happens in one of the following cases:

(1) There are two distinct monomials f_t and f_s of f such that $f(a) = f_s(a) = f_t(a)$.
(2) There exists a ghost monomial f_t of f such that $f(a) = f_t(a)$.

Remark 5.1.2. (i) If $\mathbf{a}, \mathbf{b} \in F^n$ are ghost roots of $f, g \in F[\Lambda]$ respectively, then both $\mathbf{a}$ and $\mathbf{b}$ are ghost roots of the product fg.
(ii) In view of Proposition 3.4.2, any ghost root $\mathbf{a}$ of f^k for $k \geq 1$ is a ghost root of f.

Definition 5.1.3. A set $A \subseteq \mathcal{T}^{(n)}$ is said to be a **corner locus** if A is a set of the form

$$A = \{x \in F^n \ : \ \forall f \in S,\ x \text{ is a ghost root of } f\} \tag{5.1}$$

for some $S \subset F[\Lambda]$. We write $\mathrm{C}_{\mathrm{loc}}(S)$ to denote the corner locus defined by S.

In the case where $S = \{f_1, ..., f_m\}$ is finite, we say that A is a **finitely generated** corner locus and write $A = \mathrm{C}_{\mathrm{loc}}(f_1, ..., f_m)$.

Lemma 5.1.4. *Suppose $f = f' + f'' \in F[\Lambda]$. Then $\mathbf{a} \in \mathrm{C}_{\mathrm{loc}}(f)$ iff one of the following hold:*

(i) $\mathbf{a} \in \mathrm{C}_{\mathrm{loc}}(f')$ *with* $f'(\mathbf{a}) \geq f''(\mathbf{a})$;
(ii) $\mathbf{a} \in \mathrm{C}_{\mathrm{loc}}(f'')$ *with* $f''(\mathbf{a}) \geq f'(\mathbf{a})$;
(iii) $f'(\mathbf{a}) \cong_\nu f''(\mathbf{a})$.

Proof. These are the only ways to get a ghost value. □

This concise formulation enables us to treat tropical varieties algebraically as simultaneous roots of sets of polynomials. In particular, a **hypersurface** now is viewed as the corner locus of a single polynomial. When the polynomials are tangible, the hypersurfaces are those from usual tropical geometry.

We write $\mathbf{2}^S$ for the power set of a set S.

In view of Definition 5.1.3, we define an operator $\mathrm{C}_{\mathrm{loc}} : \mathbf{2}^{F[\Lambda]} \to \mathbf{2}^{F^{(n)}}$ given by

$$\mathrm{C}_{\mathrm{loc}} : S \subset F[\Lambda] \mapsto \mathrm{C}_{\mathrm{loc}}(S). \tag{5.2}$$

We now proceed to study the behavior of the $\mathrm{C}_{\mathrm{loc}}$ operator.

Remark 5.1.5. If $S_A, S_B \subseteq F[\Lambda]$ then

$$S_A \subseteq S_B \Rightarrow \mathrm{C}_{\mathrm{loc}}(S_B) \subseteq \mathrm{C}_{\mathrm{loc}}(S_A). \tag{5.3}$$

Let $\{S_i\}_{i\in I}$ be a family of subsets of $F[\Lambda]$ for some index set I. Then $\bigcap_{i\in I} \mathrm{C}_{\mathrm{loc}}(S_i)$ is a corner locus and

$$\bigcap_{i\in I} \mathrm{C}_{\mathrm{loc}}(S_i) = \mathrm{C}_{\mathrm{loc}}\left(\bigcup_{i\in I} S_i\right) \ ; \ \bigcup_{i\in I} \mathrm{C}_{\mathrm{loc}}(S_i) \subseteq \mathrm{C}_{\mathrm{loc}}\left(\bigcap_{i\in I} S_i\right). \tag{5.4}$$

In particular, $\mathrm{C}_{\mathrm{loc}}(S) = \bigcap_{f\in S} \mathrm{C}_{\mathrm{loc}}(f)$.

Proof. First, equality (5.3) is a direct set-theoretical consequence of the definition of corner loci. In turn this implies that $\mathrm{C}_{\mathrm{loc}}(\bigcup_{i\in I} S_i) \subseteq \mathrm{C}_{\mathrm{loc}}(S_i)$ for each $i \in I$ and thus $\mathrm{C}_{\mathrm{loc}}(\bigcup_{i\in I} S_i) \subseteq \bigcap_{i\in I} \mathrm{C}_{\mathrm{loc}}(S_i)$. Conversely, if $x \in F^{(n)}$ is in $\bigcap_{i\in I} \mathrm{C}_{\mathrm{loc}}(S_i)$ then $\mathbf{a} \in \mathrm{C}_{\mathrm{loc}}(S_i)$ for every $i \in I$, which means that $\mathbf{a}$ is a common ghost root of $\{f \ : \ f \in S_i\}$. Thus $\mathbf{a}$ is a common ghost root of $\{f \ : \ f \in \bigcup_{i\in I} S_i\}$ which yields that $\mathbf{a} \in \mathrm{C}_{\mathrm{loc}}(\bigcup_{i\in I} S_i)$. For the second equation (inclusion) in (5.4), $\bigcap_{i\in I} S_i \subseteq S_j$ for each $j \in I$. Thus, by (5.3), $\mathrm{C}_{\mathrm{loc}}(S_j) \subseteq \mathrm{C}_{\mathrm{loc}}(\bigcap_{i\in I} S_i)$, and so, $\bigcup_{i\in I} \mathrm{C}_{\mathrm{loc}}(S_i) \subseteq \mathrm{C}_{\mathrm{loc}}(\bigcap_{i\in I} S_i)$. □

The next lemma deals with the case where A and B are finitely generated.

Lemma 5.1.6. *If $A = \mathrm{C}_{\mathrm{loc}}(f_1, ..., f_s)$ and $B = \mathrm{C}_{\mathrm{loc}}(g_1, ..., g_t)$, then $A \cap B = \mathrm{C}_{\mathrm{loc}}(f_1, ..., f_s, g_1, ..., g_t)$, and $A \cup B = \mathrm{C}_{\mathrm{loc}}(\{f_i g_j\}_{i=1,j=1}^{s,t})$.*

Proof. In view of Definition 5.1.3 and Remarks 5.1.2 and 5.1.5, we only need to prove that $A \cup B = \mathrm{C}_{\mathrm{loc}}\left(\{f_i g_j\}_{i=1,j=1}^{s,t}\right)$. Indeed, for each $1 \le i \le s$ and each $1 \le j \le t$, $A \subseteq \mathrm{C}_{\mathrm{loc}}(f_i)$ and $B \subseteq \mathrm{C}_{\mathrm{loc}}(g_j)$. Thus $A \cup B \subseteq \mathrm{C}_{\mathrm{loc}}(f_i) \cup \mathrm{C}_{\mathrm{loc}}(g_j) = \mathrm{C}_{\mathrm{loc}}\left(\{f_i g_j\}\right)$ and so

$$A \cup B \subseteq \bigcap_{i=1,j=1}^{s,t} \mathrm{C}_{\mathrm{loc}}\left(\{f_i g_j\}\right) = \mathrm{C}_{\mathrm{loc}}\left(\{f_i g_j\}_{i=1,j=1}^{s,t}\right).$$

On the other hand, if $a \notin A \cup B$ then there exist some i_0 and j_0 such that $a \notin \mathrm{C}_{\mathrm{loc}}(f_{i_0})$ and $a \notin \mathrm{C}_{\mathrm{loc}}(g_{j_0})$. Thus $a \notin \mathrm{C}_{\mathrm{loc}}(f_{i_0}) \cup \mathrm{C}_{\mathrm{loc}}(g_{j_0}) = \mathrm{C}_{\mathrm{loc}}(f_{i_0} g_{j_0})$. So $a \notin \bigcap_{i=1,j=1}^{s,t} \mathrm{C}_{\mathrm{loc}}\left(\{f_i g_j\}\right) = \mathrm{C}_{\mathrm{loc}}\left(\{f_i g_j\}_{i=1,j=1}^{s,t}\right)$, proving the opposite inclusion. □

In [20] the following density condition was introduced to pick out those varieties defined by several polynomials, which are more in line with the usual tropical viewpoint.

Definition 5.1.7. A corner locus Z is **admissible** if it satisfies the following property:

If f and g agree on a ν-dense subset of Z then $f = g$.

A **variety** is an admissible corner locus that is not the union of two admissible corner loci.

5.2. 1^ν-sets.

We present a method for describing a 'corner locus' by the analogous concept for ν-kernels, which we call a 1^ν-set. This in turn sets the stage for applying the theory of ν-kernels to tropical geometry, focusing on the correspondence between ν-kernels and 1^ν-sets. We introduce the geometric notion of '1^ν-set' and the algebraic notion of 'ν-kernel of a 1^ν-set', the respective analogs of affine varieties and their ideals. To obtain a Zariski-type correspondence, we define a pair of operators 1_{loc} and $\mathcal{K}ern$ where 1_{loc} maps a ν-kernel to its 1^ν-set and $\mathcal{K}ern$ maps a 1^ν-set to its corresponding ν-kernel.

Throughout this section, we take F to be a ν-semifield†. $\overline{F(\Lambda)}$ was defined in Definition 3.3.5.

Definition 5.2.1. A **kernel root** of $f \in \mathrm{Fun}(F^{(n)}, F)$ is an element $\mathbf{a} \in F^{(n)}$ such that $f(\mathbf{a}) \cong_\nu \mathbf{1}_F$.

For $S \subseteq \overline{F(\Lambda)}$, define the subset $1_{\mathrm{loc}}(S)$ of $F^{(n)}$ as

$$1_{\mathrm{loc}}(S) = \{\mathbf{a} \in F^{(n)} \ : \ f(\mathbf{a}) \cong_\nu 1, \ \forall f \in S\}. \tag{5.5}$$

We write $1_{\mathrm{loc}}(f)$ for $1_{\mathrm{loc}}(\{f\})$.

Lemma 5.2.2. *The following hold for $f, g \in \overline{F(\Lambda)}$:*

(i) $1_{\mathrm{loc}}(fg) = 1_{\mathrm{loc}}(f+g) = 1_{\mathrm{loc}}(f) \cap 1_{\mathrm{loc}}(g)$ *for all* $f, g \geq_\nu 1$.
(ii) $1_{\mathrm{loc}}(f \wedge g) = 1_{\mathrm{loc}}(f) \cup 1_{\mathrm{loc}}(g)$ *for all* $f, g \leq_\nu 1$.

Proof. (i) $(f+g)(\mathbf{a}) = f(\mathbf{a}) + g(\mathbf{a}) \cong_\nu 1$ if and only if $f(\mathbf{a}) \cong_\nu 1$ and $g(\mathbf{a}) \cong_\nu 1$ (since both $f, g \geq_\nu 1$). Analogously the same holds for $f(\mathbf{a})g(\mathbf{a}) \cong_\nu 1$.
(ii) $(f \wedge g)(\mathbf{a}) = f(\mathbf{a}) \wedge g(\mathbf{a}) \cong_\nu 1$, if and only if $f(\mathbf{a}) \cong_\nu 1$ or $g(\mathbf{a}) \cong_\nu 1$. □

Lemma 5.2.3. $1_{\mathrm{loc}}(f) = 1_{\mathrm{loc}}(f^{-1}) = 1_{\mathrm{loc}}(|f|) = 1_{\mathrm{loc}}(f \wedge f^{-1})$.

Proof. Clearly $f(\mathbf{a}) \cong_\nu 1$ iff $f^{-1}(\mathbf{a}) \cong_\nu 1$, in which case $|f|(\mathbf{a}) \cong_\nu (f \wedge f^{-1})(\mathbf{a}) \cong_\nu 1$. Conversely, if $|f|(\mathbf{a}) \cong_\nu 1$ then both $f(\mathbf{a}), f^{-1}(\mathbf{a}) \leq_\nu 1$, implying $f(\mathbf{a}), f^{-1}(\mathbf{a}) \cong_\nu 1$. The last assertion is by Lemma 5.2.2(ii). □

Definition 5.2.4. A subset $Z \subset F^{(n)}$ is said to be a **1^ν-set** if there exists a subset $S \subset \overline{F(\Lambda)}$ such that $Z = 1_{\mathrm{loc}}(S)$.

Proposition 5.2.5. *For $S_i \subset \overline{F(\Lambda)}$ the following statements hold:*

(1) $S_1 \subseteq S_2 \Rightarrow 1_{\mathrm{loc}}(S_2) \subseteq 1_{\mathrm{loc}}(S_1)$.
(2) $\bigcap_{i\in I} 1_{\mathrm{loc}}(S_i) = 1_{\mathrm{loc}}(\bigcup_{i\in I} S_i)$ *for any index set I and in particular,* $1_{\mathrm{loc}}(S) = \bigcap_{f\in S} 1_{\mathrm{loc}}(f)$.

Proof. The assertions are formal. □

Lemma 5.2.6. $1_{\mathrm{loc}}(S_1) = 1_{\mathrm{loc}}(\langle S_1 \rangle)$.

Proof. For any f, g such that $f(\mathbf{a}) \cong_\nu g(\mathbf{a}) \cong_\nu 1$,

$$(fg)(\mathbf{a}) = f(\mathbf{a})g(\mathbf{a}) \cong_\nu 1 \cdot 1 = 1, \quad f^*(\mathbf{a}) \cong_\nu f(\mathbf{a})^* \cong_\nu 1^{-1} = 1, \quad \text{and} \quad (f+g)(\mathbf{a}) = f(\mathbf{a}) + g(\mathbf{a}) \cong_\nu 1+1 \cong_\nu 1.$$

To show that convexity is preserved for $f_1, ..., f_t \in \overline{F(\Lambda)}$ with $\sum_{i=1}^t f_i \cong_\nu 1$ and for any $g_1, ..., g_t$ in $1_{\mathrm{loc}}(S_1)$, we note that

$$\left(\sum_{i=1}^t f_i g_i\right)(\mathbf{a}) = \sum_{i=1}^t f_i(\mathbf{a}) g_i(\mathbf{a}) \cong_\nu \sum_{i=1}^t (f_i(\mathbf{a}) \cdot 1) \cong_\nu \sum_{i=1}^t f_i(\mathbf{a}) \cong_\nu 1.$$

□

Proposition 5.2.7. [28, Proposition (4.2.6)] *Let K_i be ν-kernels of $\overline{F(\Lambda)}$, and let $Z_i = 1_{\mathrm{loc}}(K_i)$ be their corresponding 1^ν-sets. Then the following statements hold:*

$$1_{\mathrm{loc}}(K_1 K_2) = Z_1 \cap Z_2; \tag{5.6}$$

$$1_{\mathrm{loc}}(\cap K_i) = \cup Z_i. \tag{5.7}$$

Note 5.2.8. The proposition fails miserably when we take $\wedge$ of an infinite set. For example, we could take f_m to be a series of constants approaching 1. Then $\inf\{f_m\} = 1$, whose 1^ν-set is $F^{(n)}$, although $1_{\mathrm{loc}}(f_m) = \emptyset$ for each m.

Definition 5.2.9. A 1^ν-set Z is said to be **principal** if there exists $f \in \overline{F(\Lambda)}$ such that $Z = 1_{\mathrm{loc}}(f)$.

Definition 5.2.10. Denote the collection of 1^ν-sets in $F^{(n)}$ by 1^ν-Set$(F^{(n)})$ and the collection of principal 1^ν-sets in $F^{(n)}$ by P1^ν-Set. ($(F^{(n)})$ will be understood in the context.)

Although principal 1^ν-sets are analogous to hypersurfaces, they are more pervasive because of Lemma 5.2.2.

Proposition 5.2.11. P1^ν-Set *is closed under finite unions and intersections.*

Proof. By the lemma. □

As a special case of Proposition 5.2.7 we have

Corollary 5.2.12. *For $f, g \in \overline{F(\Lambda)}$,*

$$1_{\mathrm{loc}}(\langle f \rangle \langle g \rangle) = 1_{\mathrm{loc}}(f) \cap 1_{\mathrm{loc}}(g); \tag{5.8}$$

$$1_{\mathrm{loc}}(\langle f \rangle \cap \langle g \rangle) = 1_{\mathrm{loc}}(f) \cup 1_{\mathrm{loc}}(g). \tag{5.9}$$

Proposition 5.2.13. *If $\langle f \rangle$ is a principal ν-kernel generated by $f \in \overline{F(\Lambda)}$, then $1_{\mathrm{loc}}(f) = 1_{\mathrm{loc}}(\langle f \rangle)$. Consequently, $1_{\mathrm{loc}}(f) = 1_{\mathrm{loc}}(f')$ for any generator f' of $\langle f \rangle$.*

Proof. By Corollary 4.1.21, $g \in \overline{F(\Lambda)}$ is in $\langle f \rangle$ if and only if there exists some $n \in \mathrm{N}$ such that $|f|^{-n} \leq_\nu g \leq_\nu |f|^n$. Hence, for every $g \in \langle f \rangle$, $1 = 1^{-n} \leq_\nu g(\mathbf{a}) \leq_\nu 1^n = 1$ for each $\mathbf{a} \in 1_{\mathrm{loc}}(f)$, implying $g(\mathbf{a}) \cong_\nu 1$.
In particular, $1_{\mathrm{loc}}(f) = 1_{\mathrm{loc}}(\langle f \rangle) = 1_{\mathrm{loc}}(\langle f' \rangle) = 1_{\mathrm{loc}}(f')$. □

5.3. **Kernels of 1^ν-sets.**

In view of Lemma 5.2.6, the operator 1_{loc} can be restricted to ν-kernels. In the other direction, we now construct an operator that associates a ν-kernel of the ν-semifield† of fractions $\overline{F(\Lambda)}$ to any subset of $F^{(n)}$. Then we note that the operator $\mathcal{K}ern$ and the operator 1_{loc} defined in the previous subsection are inverses of each other.

Definition 5.3.1. Given a subset Z of $F^{(n)}$, we define

$$\mathcal{K}ern(Z) = \{f \in \overline{F(\Lambda)} \ : \ f(\mathbf{a}) \cong_\nu 1, \ \forall \mathbf{a} \in Z\}. \tag{5.10}$$

$\mathcal{K}ern(Z)$ is a ν-kernel of $\overline{F(\Lambda)}$, by the same argument as given in Lemma 5.2.6.

Remark 5.3.2. The following statements hold for $Z, Z_i \subset F^{(n)}$:

(1) If $Z_1 \subseteq Z_2$, then $\mathcal{K}ern(Z_2) \subseteq \mathcal{K}ern(Z_1)$.
(2) $\mathcal{K}ern(\bigcup_{i\in I} Z_i) = \bigcap_{i\in I} \mathcal{K}ern(Z_i)$.
(3) $K \subseteq \mathcal{K}ern(1_{\mathrm{loc}}(K))$ for any ν-kernel K of $\overline{F(\Lambda)}$.
(4) $Z \subseteq 1_{\mathrm{loc}}(\mathcal{K}ern(Z))$.

Lemma 5.3.3. *If $Z \subset F^{(n)}$, then $\mathcal{K}ern(Z)$ is a ν-kernel.*

Proof. It is closed under multiplication, and if $g_1 + g_2 \cong_\nu 1$ then $(g_1 f + g_2 f)(\mathbf{a}) = g_1(\mathbf{a}) + g_2(\mathbf{a}) \cong_\nu \mathbf{a}$. □

Definition 5.3.4. A **$\mathcal{K}$-kernel** of $\overline{F(\Lambda)}$ is a ν-kernel of the form $\mathcal{K}ern(Z)$, where Z is a 1^ν-set.

Lemma 5.3.5.

(1) $1_{\mathrm{loc}}(\mathcal{K}ern(Z)) = Z$ *for any 1^ν-set Z.*
(2) $\mathcal{K}ern(1_{\mathrm{loc}}(K)) = K$, *for any $\mathcal{K}$-kernel K.*

Proof. This is a standard argument, applying Remark 5.3.2 to the reverse inclusion of Proposition 5.2.5. □

By Lemma 5.3.5, we have

Theorem 5.3.6. *There is a $1:1$, order-reversing correspondence*

$$\{1^\nu\text{-sets of } \ F^{(n)}\} \to \{\mathcal{K} - \text{kernels of } \ \overline{F(\Lambda)}\}, \tag{5.11}$$

given by $Z \mapsto \mathcal{K}ern(Z)$; the reverse map is given by $K \mapsto 1_{\mathrm{loc}}(K)$.

One of the main goals in this paper is to find an intrinsic characterization of $\mathcal{K}$-kernels, especially the principal $\mathcal{K}$-kernels; this is only done in Corollary 8.0.14.

Definition 5.3.7. The **K-variety** of a set $S \subseteq \mathrm{Fun}(F^{(n)}, F)$, denoted 1-loc$(S)$, is $\bigcap\{\text{1-loc}(f) : f \in S\}$. (Usually S is taken to be finite.)

We need to cut down the class of K-varieties for tropical applications, in view of the following examples.

Example 5.3.8.

(1) *Let $f = \lambda_1 + \lambda_2$ and $g = \mathbf{1}_F$.*
(2) *Let $f = \lambda_1 + \lambda_2 + \mathbf{1}_F$ and $g = \lambda_1^3 + \lambda_2^2 + \mathbf{1}_F$.*
(3) *Let $f = \lambda_1 + \lambda_2 + \mathbf{1}_F$ and $g = \lambda_1^2 + \lambda_2^2 + \mathbf{1}_F$.*

In each of (1)–(3), $\text{1-loc}(\frac{f}{g})_{\mathcal{T}} = \{(\alpha, \mathbf{1}_F) : \alpha \leq_\nu \mathbf{1}_F\} \cup \{(\mathbf{1}_F, \alpha) : \alpha \leq_\nu \mathbf{1}_F\}$ is not a variety.

(4) *Let $f = \lambda_1 + \lambda_2 + \mathbf{1}_F$ and $g = \lambda_1 + \lambda_2$. Then $\text{1-loc}(\frac{f}{g})_{\mathcal{T}} = \{(\alpha, \alpha) : \alpha \geq_\nu \mathbf{1}_F\}$.*

5.4. **The coordinate ν-semifield† of a 1^ν-set.**

Definition 5.4.1. For $X \subset F^{(n)}$, The **coordinate ν-semifield†** $F(X)$ of a 1^ν-set X is the set of restrictions of the rational functions $\overline{F(\Lambda)}$ to X.

$$\phi_X \ : \ \overline{F(\Lambda)} \ \to \ F(X)$$

denotes the restriction map $h \mapsto h|_X$.

Proposition 5.4.2. *ϕ_X is an onto ν-semifield† homomorphism.*

Proof. Straightforward verification. □

Proposition 5.4.3. *$F(X)$ is a ν-domain†, isomorphic to $\overline{F(\Lambda)}/\mathcal{K}ern(X)$.*

Proof. The restriction map has ν-kernel equal to those functions f which restrict to 1, which is $\mathcal{K}ern(X)$. □

When $X' \supseteq X$ is another 1^ν-set, further restriction gives us a semiring† homomorphism $F[X] \to F[X']$, and chains of these homomorphisms give us an algebraic view of dimension, which is studied at the end of [28].

6. The transition between tropical varieties and K-varieties

In view of Lemma 5.3.3, we would like to pass back and forth from tropical varieties to K-varieties. This is one of our main themes.

6.1. The hat construction.

We start by passing to the corner kernel locus from the corner locus obtained from (super)tropical polynomials. Towards this end, we formulate the following notion. We say that a function f **dominates** g at $\mathbf{a}$ if $f(\mathbf{a}) \geq_\nu g(\mathbf{a})$; f **dominates** g if f dominates g at each point $\mathbf{a}$.

Definition 6.1.1. A **molecule** is a rational function hg^* where h is a monomial and g is a polynomial. The molecule is **tangible** if h and g are tangible. (In this case, $hg^* = \frac{h}{g}$.) Given a polynomial $f \in F[\Lambda] = \sum_{i=1}^k f_i$, written as a sum of monomials, define the **molecules of** f to be the molecules

$$\widehat{f_i} = f_i \left(\sum_{j \neq i} f_j \right)^* \in \overline{F(\Lambda)}, \tag{6.1}$$

and

$$\widehat{f} = \sum_{i=1}^k \widehat{f_i} \tag{6.2}$$

Usually f is tangible, in which case we have

$$\widehat{f_i} = \frac{f_i}{\sum_{j \neq i} f_j} \in \overline{F(\Lambda)}, \tag{6.3}$$

Definition 6.1.2. Given a polynomial $f = \sum h_i$ written as a sum of monomials, and selecting one of these monomials $h = h_i$, define $f_{(h)}$ to be $\sum_{h_j \neq h} h_j$. Then

$$\hat{f} = \sum_h h f_{(h)}^*.$$

Lemma 6.1.3. *Given $f = \sum h$ written as a sum of monomials, then $\hat{f} \in \mathcal{T}^+(\lambda_1, \ldots, \lambda_n)_\nu$.*

Proof. Write $f = \sum h$ and $\hat{f} = \sum_h h f_{(h)}^*$. If some monomial h dominates at $\mathbf{a}$, then

$$\hat{f}(\mathbf{a}) = h f_{(h)}^*(\mathbf{a}) \geq_\nu \mathbf{1}_F.$$

Thus, we may assume that two monomials g, h dominate at $\mathbf{a}$, and then

$$\hat{f}(\mathbf{a}) = g f_{(g)}^*(\mathbf{a}) + h f_{(h)}^*(\mathbf{a}) = \mathbf{1}_F^\nu.$$

□

Lemma 6.1.4. *(i) $\widehat{f_i}(\mathbf{a}) \geq_\nu \mathbf{1}_F$ iff $f_i(\mathbf{a})$ dominates $f(\mathbf{a})$.*

(ii) $\widehat{f_i}(\mathbf{a}) \cong_\nu \mathbf{1}_F$ iff, for some $j \neq i$, $f_i(\mathbf{a}) = f_j(\mathbf{a})$ which dominates each monomial. This means $f_i(\mathbf{a})$ dominates in $f(\mathbf{a})$, and $\mathbf{a} \in \mathrm{C}_{\mathrm{loc}}(f)$.

Proof. (i) Each side says that the numerator of $\widehat{f_i}$ dominates the denominator.

(ii) Clearly $\widehat{f_i}(\mathbf{a}) \cong_\nu \mathbf{1}_F$ iff, for some $j \neq i$, $f_i(\mathbf{a}) = f_j(\mathbf{a})$ which dominates each other monomial. This means $f_i(\mathbf{a}) + f_j(\mathbf{a}) = f(\mathbf{a})^\nu$, implying $f(\mathbf{a}) = f(\mathbf{a})^\nu \in \mathcal{G}$. □

Proposition 6.1.5. *The following conditions are equivalent:*

(i) $\widehat{f}(\mathbf{a}) = \mathbf{1}_F^\nu$;
(ii) $\widehat{f}(\mathbf{a}) \cong_\nu \mathbf{1}_F$;
(iii) $\mathbf{a} \in \mathrm{C}_{\mathrm{loc}}(f)$.

Proof. (i) → (ii). Obvious.

(ii) → (iii). By Lemma 6.1.4, there are two monomials f_i and f_j that dominate at $\mathbf{a}$, and so

$$\mathbf{1}_F^\nu = \widehat{f_i}(\mathbf{a}) + \widehat{f_j}(\mathbf{a}),$$

which then is $\widehat{f}(\mathbf{a})$.

(iii) → (i). $\widehat{f_i}(\mathbf{a}) = \widehat{f_j}(\mathbf{a}) \cong_\nu \mathbf{1}_F^\nu$, so

$$\mathbf{1}_F^\nu = \widehat{f_i} + (\mathbf{a}) = \widehat{f_j}(\mathbf{a}) = \widehat{f}(\mathbf{a}).$$

□

Corollary 6.1.6. $1_{\mathrm{loc}}(\widehat{f}) = \mathrm{C}_{\mathrm{loc}}(f)$.

A **principal corner locus** is a set of ghost roots of a supertropical polynomial.

Corollary 6.1.7. *Any principal corner locus is a principal 1^ν-set.*

Remark 6.1.8. For $S \subset \mathcal{F}(F[\Lambda])$, let $\widehat{S} = \{\widehat{f} \ : \ f \in S\} \subseteq \overline{F(\Lambda)}$. Then by Remark 5.1.5 and Proposition 5.2.5,

$$\mathrm{C}_{\mathrm{loc}}(S) = \bigcap_{g \in \widehat{S}} 1_{\mathrm{loc}}(g) = 1_{\mathrm{loc}}(\widehat{S}).$$

Thus, the map $\mathrm{C}_{\mathrm{loc}}(f) \mapsto 1_{\mathrm{loc}}(\widehat{f})$ extends to a map

$$\Phi : \ \mathrm{C}_{\mathrm{loc}}(F^{(n)}) \ \to \ 1^\nu\text{-Set}(F^{(n)}),$$

where $\Phi : \mathrm{C}_{\mathrm{loc}}(f) \mapsto 1^\nu\text{-Set}(\widehat{f})$. In particular, taking only finitely generated corner loci, and recalling that finite intersections and unions of principal 1^ν-sets are principal 1^ν-sets, Φ sends every finitely generated corner locus to a principal 1^ν-set.

Lemma 6.1.9. [28, Lemma (10.2.5)] *Let $f = \sum_{i=1}^k f_i \in F[\Lambda]$. Then for $1 \le i, j \le k$ such that $i \ne j$ and $\widehat{f_i}(\mathbf{a}) \cong_\nu \widehat{f_j}(\mathbf{a})$,*

$$\text{(6.4)} \qquad \textit{Either} \qquad \widehat{f_i}(\mathbf{a}) \cong_\nu \widehat{f_j}(\mathbf{a}) \cong_\nu 1 \qquad \textit{or} \qquad \widehat{f_i}(\mathbf{a}) \cong_\nu \widehat{f_j}(\mathbf{a}) \textit{ are dominated by } \widehat{f} \textit{ at } \mathbf{a}.$$

Proof. Take the two cases, where either f_i and f_j both dominate, or neither dominates. □

Lemma 6.1.10. *For $f, g \in F[\Lambda]$ with g tangible, over a ν-semifield† F (where $\Lambda = \{\lambda_1, \dots, \lambda_n\}$), the union of $1\text{-loc}(\frac{f}{g})$ with the corner loci of f and g, is the corner locus of $f + g$.*

Proof. $\frac{f}{g}(\mathbf{a}) \cong_\nu \mathbf{1}_F$ iff $f(\mathbf{a}) \cong_\nu g(\mathbf{a})$, which happens when $f(\mathbf{a}) + g(\mathbf{a}) \in \mathcal{G}$. □

Proposition 6.1.11. *An element $\mathbf{a} \in F^{(n)}$ is a ν-kernel root of $\hat{f}$ iff $\mathbf{a}$ is a ghost root of f.*

Proof. Apply Lemma 6.1.10 taking $g = f_{(h)}$. □

We can state this more explicitly.

Proposition 6.1.12. *Suppose $f = \sum_{i=1}^t h_i \in \overline{F[\lambda_1, \dots, \lambda_n]}$, written as a sum of monomials. Then*

$$\hat{f} = f^t \left(\prod_{i=1}^t f_{h_i} \right)^*.$$

Proof. (i) Since $\hat{f} = \sum_i h_i f_{h_i}^* = (\sum_i h_i \prod_{j \ne i} f_{h_j})(\prod_j f_{h_j})^*$, it suffices to prove that $\sum_i h_i \prod_{j \ne i} f_{h_j} = f^t$. But each h_i^t appears in both sides, and every other product of length t of the h_j appears as a ghost in both sides. □

We get Proposition 6.1.11 as a consequence, since the kernel roots must be precisely those $\mathbf{a}$ for which each $\prod f_{h_i}(\mathbf{a}) \cong_\nu f(\mathbf{a})$, which are the ghost roots.

Remark 6.1.13. Applying Proposition 6.1.12 to $\hat{f}_{(h_i)}$ and observing that f is dominated at each ghost root in two monomials, by definition, we see that the 1^ν-set of $\hat{f}_{(h_i)}$ is precisely the 1^ν-set of $f_{(h_i)}^{t-1}(\prod_{j \ne i} f_{h_j})^*$. From this point of view, we can cut down one summand when passing to 1^ν-sets.

Remark 6.1.14. The hat construction in Perri's dissertation [28] is related to the bend congruences of [9, 26, 25] (up to ν-equivalence), although with the proviso that these authors deal with polynomials per se, whereas at the outset we view polynomials as functions. (This gap is seen to be bridged by region fractions, which produce **region ν-kernels**, cf. Definition 9.2.7, enabling us to identify those ν-kernels which are trivial in a given neighborhood, thereby indicating which local information may be lost, as explained in the decomposition of Theorem 9.4.8.)

Namely, take a polynomial $f = \sum_i f_i = h + f_{(h)}$ where $h = f_j$ is one of the f_i dominating at $\mathbf{a}$ and $f_{(h)} := \sum_{i\neq j} f_i$; the **bend congruence** is defined to be the congruence in $F[\Lambda]$ generated by $(f, f_{(h)})$. But when viewed as a function, this could be matched with the principal ν-kernel defined by the rational function

$$\frac{f}{f_{(h)}} = \sum_i \frac{f_i}{f_{(h)}} \cong_\nu \frac{h}{f_{(h)}} \cong_\nu \hat{f}_j,$$

and we can reverse directions.

6.2. Examples of 1^ν-sets.

Example 6.2.1. In these examples, we write $\mathbf{a} = (a_1, \dots, a_n)$.

(i) Take the tropical line $f = \lambda_1 + \lambda_2 + 1$. Its corresponding 1^ν-set is defined by the rational function

$$\widehat{f} = \frac{\lambda_1}{\lambda_2+1} + \frac{\lambda_2}{\lambda_1+1} + \frac{1}{\lambda_1+\lambda_2} = \frac{(\lambda_1+\lambda_2+1)^3}{(\lambda_1+1)(\lambda_2+1)(\lambda_1+\lambda_2)}. \tag{6.5}$$

Moreover, in Remark 6.1.13 it is shown that any of the three terms above can be omitted, and for example we could use instead

$$\frac{\lambda_1}{\lambda_2+1} + \frac{\lambda_2}{\lambda_1+1} = \frac{\lambda_1^2+\lambda_2^2+\lambda_1+\lambda_2}{(\lambda_1+1)(\lambda_2+1)}. \tag{6.6}$$

Although these two functions (6.5) and (6.6) define the same 1^ν-set, they define different ν-kernels, since $\lambda_1+\lambda_2$ is not in the ν-kernel generated by $\frac{\lambda_1}{\lambda_2+1} + \frac{\lambda_2}{\lambda_1+1}$. Indeed, if we take the point $\mathbf{a} = (\alpha, \beta)$ for $\alpha < \beta < 1$, we get $\frac{1}{\alpha}$ in (6.5) but β in (6.6), and their ratio can be whatever we want, so the condition of Proposition 4.1.18 fails. This ambiguity motivates much of the theory espoused later in this paper.

(ii) The ghost roots of $a_2 \leq_\nu a_1 \cong_\nu 1$ and $a_1 \leq_\nu a_2 \cong_\nu 1$ provide two rays, but we are missing the ray $1 \leq_\nu a_1 \cong_\nu a_2$.

(iii) The 1^ν-set of the rational function $\frac{1}{\lambda_1+\lambda_2} + \frac{\lambda_1}{\lambda_2} + \frac{\lambda_2}{\lambda_1}$ is comprised only of the ray $1 \leq_\nu a_1 \cong_\nu a_2$.

(iv) Next, we take the tropical line $g = \lambda_1 + \lambda_2 + 2$. Its corresponding 1^ν-set is defined by the rational function $\widehat{g} = \frac{\lambda_1}{\lambda_2+2} + \frac{\lambda_2}{\lambda_1+2} + \frac{2}{\lambda_1+\lambda_2}$.

(v) We take the corner locus defined by the common ghost roots of f and g, which correspond to the 1^ν-set defined by the rational function $\widehat{f} + \widehat{g}$. Taking the three dominant monomials yields $\frac{\lambda_1}{\lambda_2+1} + \frac{\lambda_2}{\lambda_1+1} + \frac{2}{\lambda_1+\lambda_2}$, which again cannot come from a single polynomial.

6.3. From the ν-kernel locus to hypersurfaces.

The other direction is more straightforward.

Definition 6.3.1. Given $f = \frac{h}{g} \in \overline{F(\Lambda)}$ for $g, h \in \overline{F[\lambda_1, \dots, \lambda_n]}$, define $\underline{f} = g + h$.

Remark 6.3.2. Conceivably, one could have $f = \frac{h}{g} = \frac{h'}{g'} \in \mathcal{T}^+(\lambda_1, \dots, \lambda_n)_\nu$, so the definition technically depends on how we define the representation. But in this case, $hg' = gh'$. Thus, any ghost root $\mathbf{a}$ of h is a ghost root of g or h', so we get the same ghost roots except when both the numerator and denominator have a common ghost root.

Theorem 6.3.3. *Suppose $f \in \overline{F[\lambda_1, \dots, \lambda_n]}$. Then*

$$\mathrm{C}_{\mathrm{loc}}(\underline{\hat{f}}) = 1_{\mathrm{loc}}(\hat{f}) = \mathrm{C}_{\mathrm{loc}}(f).$$

Proof. Using Proposition 6.1.12 and its proof, we can compute $\underline{\widehat{f}}$ explicitly, as $f^t + \prod_{i=1}^t f_{h_i} = f^t$ (since each $f_{h_i} \leq_\nu f$, and $\prod_{i=1}^t f_{h_i}$ does not contain any h_i^t, which are the only tangible monomials of f^t). □

In this way, we can pass from the corner locus of a single polynomial to a 1-set, and back. This procedure breaks down for tropical varieties defined by more than one polynomial, as seen in Example 6.2.1(v).

Going the other way is trickier. Suppose $f = \frac{h}{g}$. Then $\widehat{(\underline{f})} = \widehat{g+h}$.

Example 6.3.4. *Take $f = \frac{h}{g} \in \overline{F(\Lambda)}$, for $g = \lambda_1$ and $h = 2\lambda_1 + 4$. Then $\underline{f} = h$ has the ghost root 2, which is not in the 1-set of fS.*

7. Distinguishing the canonical tropical varieties

By "canonical" tropical variety we mean the kind of variety that arises in the usual study of tropical geometry, cf. [24], although there really is no truly "canonical" definition. One possibility is the definition of variety given above. Our first objective is to indicate how to obtain these via ν-kernels. In this section, following Perri's dissertation, [28], we present a method for describing a canonical tropical variety in terms of a single rational function, thereby enabling us to pass from corner loci to 1^ν-sets which often are principal. This sets the path for applying the study conducted in the previous sections to tropical geometry. We begin our discussion by introducing the notion of "corner internal" 1^ν-sets, which characterize the 1^ν-sets and ν-kernels arising from tropical varieties. Then we bring in "regular" rational functions, which avoid degeneracy in the variety.

7.1. **Corner internal functions.**

First we want to know which ν-kernels come from corner loci. Example 6.3.4 shows that we can have $f = \frac{h}{g}$, where $\mathbf{a}$ is a ghost root of h and thus of $\underline{f} = h + g$, but having $h(\mathbf{a}) >_\nu g(\mathbf{a})$, and thus $\mathbf{a}$ is not in the 1-set of f. Our next definition is designed to exclude this possibility.

Definition 7.1.1. A rational function $f \in \overline{F(\Lambda)}$ is **corner internal** if we can write $f = \frac{h}{g}$ for polynomials g, h, with $g \in \overline{\mathcal{T}[\Lambda]}$, such that every ghost root of $g + h$ is a kernel root of f.

In other words, if $g(\mathbf{a}) \leq_\nu h(\mathbf{a}) \in \mathcal{G}$ or $h(\mathbf{a}) \leq_\nu g(\mathbf{a}) \in \mathcal{G}$, then $g(\mathbf{a}) \cong_\nu h(\mathbf{a})$.

Remark 7.1.2. Suppose $f = hg^*$. Any ghost root of $g + h$ is either a ghost root of h, or of g, or else satisfies $g(\mathbf{a}) \cong_\nu h(\mathbf{a})$ in which case it is automatically in $1_{\mathrm{loc}}(f)$. Thus, corner internality is equivalent to showing that $\mathrm{C}_{\mathrm{loc}}(g) \cup \mathrm{C}_{\mathrm{loc}}(h) \subseteq 1_{\mathrm{loc}}(f)$.

Remark 7.1.3. When considering whether $f \in \overline{F(\Lambda)}$ is corner internal, the choice of g and h is important, which we call the **canonical** way of writing f. For example, $\lambda + 1 \in F(\lambda)$ is trivially corner internal, although $\frac{(\lambda+\alpha)(\lambda+1)}{\lambda+\alpha}$ for $\alpha >_\nu 1$ is not - since substituting α for λ we get $\frac{\alpha^2+\alpha^2}{\alpha+\alpha} = \frac{(\alpha^\nu)2}{\alpha^\nu)}$. Thus α is a ghost root of the numerator, which surpasses the denominator since $(\alpha^2)^\nu >_\nu \alpha^\nu$.

Also note that $g + h$ must be tangible, since otherwise some open set around $\mathbf{a}$ is in $\mathrm{C}_{\mathrm{loc}}(g + h)$ but not in $1_{\mathrm{loc}}(f)$. Thus, we write $f = \frac{h}{g}$ instead of $f = hg^*$. In particular, f must be reduced as a fraction.

Even though the canonical way of writing $f = \frac{h}{g}$ may not be unique, the following observation is enough for our purposes.

Lemma 7.1.4. *Suppose $f = \frac{h}{g} = \frac{h'}{g'}$ are two canonical ways of writing a corner internal rational function f. Then $\mathrm{C}_{\mathrm{loc}}(h + g) = \mathrm{C}_{\mathrm{loc}}(h' + g')$.*

Proof. If $\mathbf{a} \in \mathrm{C}_{\mathrm{loc}}(h+g)$, then either $g(\mathbf{a}) \cong_\nu h(\mathbf{a})$ and we are done, or say $\mathbf{a} \in \mathrm{C}_{\mathrm{loc}}(g)$. But then $h(\mathbf{a}) \leq_\nu g(\mathbf{a})$, implying $h'(\mathbf{a}) \leq_\nu g'(\mathbf{a})$, and also $\mathbf{a} \in \mathrm{C}_{\mathrm{loc}}(hg')$ since $hg' = h'g$. We are done if $\mathbf{a} \in \mathrm{C}_{\mathrm{loc}}(g')$, so we may assume that $\mathbf{a} \in \mathrm{C}_{\mathrm{loc}}(h)$. Hence $h(\mathbf{a}) \leq_\nu g(\mathbf{a}) \in \mathcal{G}$, implying $g(\mathbf{a}) \cong_\nu h(\mathbf{a})$, and again we are done. □

Thus, $\mathrm{C}_{\mathrm{loc}}(\underline{f})$ does not depend on the canonical way we write the corner internal rational function f.

Proposition 7.1.5. *A rational function f is corner internal iff $\mathrm{C}_{\mathrm{loc}}(\underline{f}) = 1_{\mathrm{loc}}(f)$, for any canonical way of writing $f = \frac{h}{g}$.*

Proof. ($\Leftarrow$) is by definition.

($\Rightarrow$) Any kernel root $\mathbf{a}$ must satisfy $g(\mathbf{a}) \cong_\nu h(\mathbf{a})$, implying $\mathbf{a}$ is a ghost root of $g + h$. □

Lemma 7.1.6. *For any polynomial $f \in F[\Lambda]$, the function $\hat{f}$ is corner internal.*

Proof. $\mathrm{C}_{\mathrm{loc}}(\underline{\hat{f}}) = 1_{\mathrm{loc}}(\hat{f})$ in view of Theorem 6.3.3. □

Theorem 7.1.7. *The correspondences $f \mapsto \hat{f}$ and $h \mapsto \underline{h}$ induce a 1:1 correspondence between hypersurfaces and 1^ν-sets of corner internal rational functions.*

Proof. Combine Proposition 7.1.5 and Lemma 7.1.6. □

Having established the importance of being corner internal, let us delve deeper into the elementary properties.

By symmetry, a rational function $f \in \overline{F(\Lambda)}$ is corner internal if and only if f^{-1} is corner internal.

Remark 7.1.8. If $f = \frac{h}{g},\ f' = \frac{h'}{g'} \in \overline{F(\Lambda)}$, then

$$f + f' = \frac{hg' + gh'}{gg'},$$

and

$$|f| = f + f^{-1} = (h^2 + g^2)(gh)^*.$$

By Proposition 3.4.2, any ghost root of $h^2 + g^2$ or of $h^2 + g^2 + gh \cong_\nu (h + g)^2$ is a ghost root of $g + h$.

Proposition 7.1.9. *A rational function $f \in \overline{F(\Lambda)}$ is corner internal if and only if $|f|$ is corner internal.*

Proof. $|f|$ and f have the same 1^ν-sets.

($\Rightarrow$) Write $f = \frac{h}{g}$ canonically. Then $|f| = hg^* + gh^* = (h^2 + g^2)(gh)^*$. By Remark 7.1.8, any ghost root $\mathbf{a}$ of $g^2 + h^2$ is a ghost root of $g + h$ and thus a kernel root of f, by definition, or else $g(\mathbf{a}) = h(\mathbf{a})$, yielding a kernel root of f. But then $\mathbf{a} \in 1_{\mathrm{loc}}(|f|)$.

($\Leftarrow$) If $\mathbf{a}$ is a ghost root of $g + h$, then, again by Remark 7.1.8, $\mathbf{a}$ is a ghost root of $g^2 + h^2 + gh$ and thus by hypothesis is in $1_{\mathrm{loc}}(|f|) = 1_{\mathrm{loc}}(f)$. □

Lemma 7.1.10. *If $f \in \overline{F(\Lambda)}$ is corner internal, then f^k, for any $k \in \mathbb{N}$ and $\sum_{i=1}^m f^{d(i)}$ with $d(1) < \cdots < d(m)$ in $\mathbb{N}$, also are corner internal.*

Proof. Everything follows easily from Remark 7.1.8. □

We recall $f \wedge g$ from (3.2) after Example 3.3.1.

Remark 7.1.11. For intuition, in view of Remark 2.5.2, $(f \wedge g)(\mathbf{a}) \cong_\nu \min\{f(\mathbf{a})^\nu, g(\mathbf{a})^\nu\}$.

Proposition 7.1.12. *If $f, f' \in \overline{F(\Lambda)}$ are corner internal, then $|f| \wedge |f'|$ is corner internal.*

Proof. Since f and f' are corner internal, $|f|$ and $|f'|$ also are corner internal. Write $|f| = \frac{h}{g}$ and $|f'| = \frac{h'}{g'}$ in the canonical ways. Now,

$$1 \leq_\nu |f| \wedge |f'| = (gh^* + g'(h')^*)^* = hh'(gh' + g'h)^*, \tag{7.1}$$

so

$$h(\mathbf{a})h'(\mathbf{a}) \geq_\nu (gh' + g'h)(\mathbf{a}). \tag{7.2}$$

We need to show that any ghost root $\mathbf{a}$ of $hh' + gh' + g'h$ satisfies

$$(hh' + gh' + g'h)(\mathbf{a}) \cong_\nu h(\mathbf{a})h'(\mathbf{a}).$$

$\leq_\nu$ follows from (7.2), so we need to show $\geq_\nu$. Clearly $\mathbf{a}$ is a ghost root of $hh' + gh'$, $hh' + g'h$, or $gh' + g'h$. In the latter case, $h(\mathbf{a})h'(\mathbf{a}) \leq_{\mathbf{a}} (gh' + g'h)(\mathbf{a})$, and we are done. Thus by symmetry we may assume that $\mathbf{a} \in \mathrm{C}_{\mathrm{loc}}(hh' + gh') = \mathrm{C}_{\mathrm{loc}}((g + h)h')$. If $\mathbf{a} \in \mathrm{C}_{\mathrm{loc}}(g + h)$ then $\mathbf{a} \in 1_{\mathrm{loc}}(|f|)$ and thus in $1_{\mathrm{loc}}(|f| \wedge |f'|)$. So we are done unless $\mathbf{a} \in \mathrm{C}_{\mathrm{loc}}(h')$. If $h'(\mathbf{a}) \geq_\nu g'(\mathbf{a})$, then $\mathbf{a} \in \mathrm{C}_{\mathrm{loc}}(g' + h') = 1_{\mathrm{loc}}(f)$ and we are done. Thus we may assume that $h'(\mathbf{a}) \leq_\nu g'(\mathbf{a})$, implying $g'(\mathbf{a})h(\mathbf{a}) \geq_\nu h'(\mathbf{a})h(\mathbf{a})$, and again we are done. □

Corollary 7.1.13. *Let $f \in \overline{F(\Lambda)}$. Then f is corner internal if and only if $|f| \wedge |\alpha|$ is corner internal for any $\alpha \neq 1$ in F.*

Proof. Write $f = \frac{h}{g}$. Since $|\alpha| > 1$, the ghost roots of $|\alpha|h + |\alpha|g + h$ and of $|\alpha|(g + h)$ are the same, and thus the kernel roots of $|f|$ and of $|f| \wedge |\alpha|$ are the same. □

Proposition 7.1.14. *There is a $1:1$ correspondence between principal corner internal 1^ν-sets and principal corner-loci.*

Proof. Take $f, f' \in \overline{F(\Lambda)}$ corner internal.

If $1_{\mathrm{loc}}(f') = 1_{\mathrm{loc}}(f)$, then

$$1_{\mathrm{loc}}(f') = 1_{\mathrm{loc}}(f) = \mathrm{C}_{\mathrm{loc}}(\underline{f}) = 1_{\mathrm{loc}}(\widehat{\underline{f}}).$$

Conversely, if $\mathrm{C}_{\mathrm{loc}}(f') = \mathrm{C}_{\mathrm{loc}}(f)$, then

$$\mathrm{C}_{\mathrm{loc}}(f') = \mathrm{C}_{\mathrm{loc}}(f) = 1_{\mathrm{loc}}(\hat{f}) = \mathrm{C}_{\mathrm{loc}}(\underline{\hat{f}}),$$

by Theorem 6.3.3. □

7.2. **Corner internal ν-kernels.**

Definition 7.2.1. A principal ν-kernel K of $\overline{F(\Lambda)}$ is said to be **corner internal** if it has a corner internal generator. In this case, the 1^ν-set $1_{\mathrm{loc}}(K)$ corresponding to K is said to be a **corner internal** 1^ν-set.

Corollary 7.2.2. *Any finite intersection of principal corner internal ν-kernels is a principal corner internal ν-kernel.*

Proof. By induction, it is enough to show that if K and K' are principal corner internal ν-kernels, then so is $K \cap K'$. Write $K = \langle f \rangle$ and $K' = \langle f' \rangle$. By Proposition 7.1.9 we may assume that $f, f' \geq 1$. By Proposition 7.1.12, $f \wedge f'$ is corner internal, which generates $K \cap K'$. □

Remark 7.2.3. If $f, g \in \overline{F(\Lambda)}$ are corner internal then $|f| + |g|$ need not be corner internal. Thus the collection of corner internal ν-kernels is not a lattice. In our study we thus take the lattice **generated** by principal corner internal ν-kernels. These elements will be shown to correspond to finitely generated corner loci that are not necessarily principal.

7.3. **The hat-construction for corner internal ν-kernels.**

Although the $f \mapsto \hat{f}$ correspondence given above yields a fast and effective correspondence from corner loci of polynomials to 1^ν-sets, it does not work so well on arbitrary rational functions $f \in \overline{F(\Lambda)}$, so we turn to a subtler but more thorough correspondence, which "explains" what makes a ν-kernel corner internal. We start with the special case of a polynomial $f = \sum_i f_i \in F[\Lambda]$, written as a sum of monomials f_i.

Definition 7.3.1. Given $f = \sum_i f_i \in F[\Lambda]$, define the rational function $\tilde{f} = \bigwedge_{i=1}^k |\widehat{f_i}|$.

Proposition 7.3.2. $\mathrm{C}_{\mathrm{loc}}(f) = 1_{\mathrm{loc}}(\tilde{f})$.

Proof. We have seen in Theorem 6.3.3 that $\mathbf{a} \in \mathrm{C}_{\mathrm{loc}}(f)$ iff $\widehat{f_i}(\mathbf{a}) = 1$ for some i. But each $|\widehat{f_i}|(\mathbf{a}) \geq 1$, so we conclude with Lemma 5.2.3 □

Thus, we have an alternative approach to that of Theorem 6.3.3, motivating some of the intricate computations we are about to make.

Theorem 7.3.3. *Suppose $f = hg^* \in \mathcal{T}(\Lambda)$ is a rational function, where $h = \sum_i h_i$ and $g = \sum_j g_j$ written as sums of monomials. Then f is corner internal if and only if the following conditions hold:*

$$\left(\bigcup_{i=1}^k 1_{\mathrm{loc}} \left(\widehat{h_i} \right) \right) \cap 1_{\mathrm{loc}}(f^* + 1) \subseteq 1_{\mathrm{loc}}(f) \tag{7.3}$$

$$\left(\bigcup_{j=1}^m 1_{\mathrm{loc}} \left(\widehat{g_j} \right) \right) \cap 1_{\mathrm{loc}}(f + 1) \subseteq 1_{\mathrm{loc}}(f). \tag{7.4}$$

Proof. Note that $f + 1 = (g + h)g^*$, whereas $f^* + 1 = (g + h)h^*$. In view of Remark 7.1.2, we need to show that $\mathrm{C}_{\mathrm{loc}}(h) \subseteq 1_{\mathrm{loc}}(f)$ is equivalent to (7.3), since the other condition is symmetric (with respect to exchanging g and h). By definition of corner internality, any ghost root $\mathbf{a}$ of h must satisfy $h(\mathbf{a}) \leq_\nu g(\mathbf{a})$, implying $\mathbf{a} \in 1_{\mathrm{loc}}(f + 1)$. Thus, checking this at each ghost root, we have $1_{\mathrm{loc}} \left(\widehat{h_i} \right) \subseteq 1_{\mathrm{loc}}(f + 1)$ for each i, or equivalently,

$$\langle f + 1 \rangle \subseteq \left\langle \widehat{h_i} \right\rangle. \tag{7.5}$$

The steps are reversible.

Now, intersecting both sides with $1_{\mathrm{loc}}(f^* + 1)$ yields

$$1_{\mathrm{loc}} \left(\widehat{h_i} \right) \cap 1_{\mathrm{loc}}(f^* + 1) \subseteq 1_{\mathrm{loc}}(f + 1) \cap 1_{\mathrm{loc}}(f^* + 1).$$

Note that this step also is reversible since $1_{\mathrm{loc}}(f^* + 1) \cup 1_{\mathrm{loc}}(f + 1) = F^{(n)}$.

Passing again from K-varieties to ν-kernels yields

$$\langle f \rangle \subseteq \left\langle \widehat{h_i} \right\rangle \langle f^* + 1 \rangle.$$

Moreover, since the above inclusion holds for every $i \in \{1, ..., k\}$ we conclude that

$$\left(\bigcup_{i=1}^k 1_{\mathrm{loc}} \left(\widehat{h_i} \right) \right) \cap 1_{\mathrm{loc}}(f^* + 1) \subseteq 1_{\mathrm{loc}}(f),$$

as desired.

□

Remark 7.3.4. Note that

$$\bigcup_{i=1}^{k} 1_{\mathrm{loc}}\left(\widehat{h_i}\right) = 1_{\mathrm{loc}}\left(\bigwedge_{i=1}^{k}\left|\widehat{h_i}\right|\right) = 1_{\mathrm{loc}}(\tilde{h})$$

and similarly

$$\bigcup_{j=1}^{m} 1_{\mathrm{loc}}\left(\widehat{g_j}\right) = 1_{\mathrm{loc}}(\tilde{g}).$$

Thus we can rewrite (7.3) and (7.4) as

$$1_{\mathrm{loc}}\left(\tilde{h}\right) \cap 1_{\mathrm{loc}}(f^*+1) \subseteq 1_{\mathrm{loc}}(f) \qquad \text{and} \qquad 1_{\mathrm{loc}}\left(\tilde{g}\right) \cap 1_{\mathrm{loc}}(f+1) \subseteq 1_{\mathrm{loc}}(f),$$

or as

$$1_{\mathrm{loc}}\left(\widehat{h}\right) \cap 1_{\mathrm{loc}}(f^*+1) \subseteq 1_{\mathrm{loc}}(f) \qquad \text{and} \qquad 1_{\mathrm{loc}}\left(\widehat{g}\right) \cap 1_{\mathrm{loc}}(f+1) \subseteq 1_{\mathrm{loc}}(f).$$

We conclude that $\widehat{f}$ is corner internal for any $f \in \overline{F(\Lambda)}$.

Remark 7.3.5. In view of Theorem 7.3.3, given $f = \frac{h}{g} \in \overline{F(\Lambda)}$, in order to obtain a corner internal rational function whose K-hypersurface contains $1_{\mathrm{loc}}(f)$ one must adjoin both

$$\left(\bigcup_{i=1}^{k} 1_{\mathrm{loc}}\left(\widehat{h_i}\right)\right) \cap 1_{\mathrm{loc}}(f^*+1)$$

and

$$\left(\bigcup_{j=1}^{m} 1_{\mathrm{loc}}\left(\widehat{g_j}\right)\right) \cap 1_{\mathrm{loc}}(f+1)$$

to the K-hypersurface of f.

writing

Definition 7.3.6. Define the map $\Phi_{CI} : \overline{F(\Lambda)} \to \overline{F(\Lambda)}$ by

$$\Phi_{CI}(f) = |f| \wedge \left(|f^*+1| + \tilde{h}\right) \wedge \left(|f+1| + \tilde{g}\right), \tag{7.6}$$

where $f = \frac{h}{g}$.

Remark 7.3.7. $\Phi_{CI}(f) = \left(|f^*+1| + \left(|f| \wedge \tilde{h}\right)\right) \wedge \left(|f+1| + \left(|f| \wedge \tilde{g}\right)\right)$, since $|f^*+1|, |f+1| \leq_\nu |f|$. This is the fraction whose K-hypersurface is formed by adjoining all the necessary points to $1_{\mathrm{loc}}(f)$ to obtain corner internality. Hence,

$$\langle \Phi_{CI}(f) \rangle = \left(\langle f^*+1\rangle \left(\langle f\rangle \cap \bigcap_{i=1}^{k} \left\langle \widehat{h_i} \right\rangle\right)\right) \cap \left(\left(\langle f+1\rangle \left(\langle f\rangle \cap \bigcap_{j=1}^{m} \langle \widehat{g_j}\rangle\right)\right)\right).$$

Corollary 7.3.8. *Let $f = \frac{h}{g} \in \overline{F(\Lambda)}$ be a rational function, where $h = \sum_{i=1}^{k} h_i$ and $g = \sum_{j=1}^{m} g_j$ are written as sums of monomials in $F[x_1, ..., x_n]$. Then*

$$1_{\mathrm{loc}}(\widehat{h+g}) = 1_{\mathrm{loc}}(f) \cup \left(\left(\bigcup_{i=1}^{k} 1_{\mathrm{loc}}\left(\widehat{h_i}\right)\right) \cap 1_{\mathrm{loc}}(f^*+1)\right) \cup \left(\left(\bigcup_{j=1}^{m} 1_{\mathrm{loc}}\left(\widehat{g_j}\right)\right) \cap 1_{\mathrm{loc}}(f+1)\right).$$

Thus $1_{\mathrm{loc}}(\widehat{h+g}) = 1_{\mathrm{loc}}(\Phi_{CI}(f))$.

Proof. Apply Lemma 7.1.2 and Remark 7.3.4 to the theorem, to get the three parts. □

Theorem 7.3.9. *If $f \in \overline{F(\Lambda)}$, then $\langle \Phi_{CI}(f) \rangle$ is corner internal, and $1_{\mathrm{loc}}(\Phi_{CI}(f)) \supseteq 1_{\mathrm{loc}}(f)$. Furthermore, $1_{\mathrm{loc}}(\Phi_{CI}(f)) = 1_{\mathrm{loc}}(f)$ if and only if f is corner internal.*

Proof. The first claim follows from Corollary 7.3.8, where $\widehat{f}$ is corner internal by Theorem 6.3.3. The second claim is straightforward from Remark 7.3.7 since

$$1_{\mathrm{loc}}\left(|f| \wedge \left(|f^* + 1| + \tilde{h}\right) \wedge (|f+1| + \tilde{g})\right) = 1_{\mathrm{loc}}(f) \cup 1_{\mathrm{loc}}\left(|f^* + 1| + \tilde{h}\right) \cup 1_{\mathrm{loc}}\left(|f+1| + \tilde{g}\right).$$

The last statement follows Theorem 7.3.3. □

Proposition 7.3.10. *For any $f = \frac{h}{g} \in \overline{F(\Lambda)}$,*

$$\langle \Phi_{CI}(\sum_{i=1}^{k} f^{d(i)}) \rangle = \langle \Phi_{CI}(f) \rangle \tag{7.7}$$

whenever $0 < d(1) < \cdots < d(k)$, and

$$\langle \Phi_{CI}(f^k) \rangle = \langle \Phi_{CI}(f) \rangle \tag{7.8}$$

for any $k \in \mathbb{Z} \setminus \{0\}$.

Proof. By the Frobenius property, $\sum_{i=1}^{k} f^{d(i)} = \frac{\sum_{i=1}^{k} h^{s+d(i)} g^{t-d(i)}}{h^s g^t} = \frac{h^{s+t} + g^{s+t}}{h^s g^t}$ where $t = |d(k)|$ and $s = |d(1)|$. So

$$\Phi_{CI}\left(\sum_{i=1}^{k} f^{d(i)}\right) = \Phi_{CI}\left(\frac{h^{s+t} + g^{s+t}}{h^s g^t}\right) = h^{s+t} + \widehat{g^{s+t}} + h^s g^t = h^{s+\widehat{t} + g^{s+t}} = (\widehat{h+g})^{s+t}.$$

(The hats are over the enitre expressions.) But $\langle (\widehat{h+g})^{s+t} \rangle = \langle \widehat{h+g} \rangle = \langle \Phi_{CI}(f) \rangle$, since $\Phi_{CI}(\frac{h^k}{g^k}) = \widehat{h^k + g^k} = (\widehat{h+g})^k$. Thus, $\langle (\widehat{h+g})^k \rangle = \langle \widehat{h+g} \rangle = \langle \Phi_{CI}(f) \rangle$. □

Corollary 7.3.11. *Let $f \in \overline{F(\Lambda)}$ be such that $f = u_1 \wedge \cdots \wedge u_k$ where $u_1, ..., u_k \in \overline{F(\Lambda)}$ are corner internal. Then f is corner internal and*

$$1_{\mathrm{loc}}(\Phi_{CI}(f)) = 1_{\mathrm{loc}}(\Phi_{CI}(u_1) \wedge \cdots \wedge \Phi_{CI}(u_k)). \tag{7.9}$$

Proof. f is corner internal, by Corollary 7.2.2. Replacing f by $|f|$, we may we assume that $f \geq 1$ (and thus each $u_i \geq 1$. Since u_i is corner internal, $1_{\mathrm{loc}}(\Phi_{CI}(u_i)) = 1_{\mathrm{loc}}(u_i)$ for $i = 1, ..., k$, and thus

$$1_{\mathrm{loc}}(\Phi_{CI}(u_1) \wedge \cdots \wedge \Phi_{CI}(u_k)) = 1_{\mathrm{loc}}(u_1 \wedge \cdots \wedge u_k) = 1_{\mathrm{loc}}(f).$$

(7.9) holds since f is corner internal. □

7.4. **Regularity.**

In the standard tropical theory, varieties have the property that their complement is dense, whereas in the supertropical theory the root set say of $\lambda^2 + 2^\nu \lambda + 3$ contains the closed interval $[1, 2]$. We would like to handle this issue through ν-kernels and their 1^ν-sets.

The relation $f(\mathbf{a}) = \frac{\sum h_i}{\sum g_j}(\mathbf{a}) \cong_\nu 1$ is studied locally in the following sense: For any $\mathbf{a} \in F^{(n)}$, there is at least one monomial h_{i0} of the numerator and at least one monomial g_{j0} of the denominator which are dominant at $\mathbf{a}$. If more than one monomial at $\mathbf{a}$ is dominant, say $\{h_{ik}\}_{k=1}^{s}$ and $\{g_{jm}\}_{m=1}^{t}$, then we have additional relations of the form $h_{i0} = h_{ik}$ and $g_{j0} = g_{jm}$.

Our motivating example: The relation $\mathbf{1}_F + \lambda = \mathbf{1}_F$ holds for all $\mathbf{a} \leq_\nu \mathbf{1}_F$. This will be an example of an **order relation**.

In general, the dominant monomials of both numerator and denominator at some point define relations $\{h_{ik} = g_{jm} \ : \ 0 \leq k \leq s, \ 0 \leq m \leq t\}$ on regions of $F^{(n)}$. Every such relation can be converted by multiplying by inverses of monomials to obtain a relation of the form $1 = \phi(\Lambda)$ with $\phi \in \overline{F(\Lambda)} \setminus F$ a Laurent monomial, and thus reduces the dimension. Note that in the case in which h_{i0} and g_{j0} singly dominate and are the same monomial, no extra relation is imposed on the region described above, so we are left only with the order relations defining the region.

In this way, two types of principal 1^ν-sets emerge from two distinct types of ν-kernels, characterized by their generators, distinguished via the following definition.

Definition 7.4.1. A rational function $f = \frac{h}{g}$ is **regular** at a point $\mathbf{a}$ in $1_{\mathrm{loc}}(f)$ if each ν-neighborhood of $\mathbf{a}$ contains a point on which h and g do not agree. Otherwise f is **irregular** at $\mathbf{a}$. f is **regular** at a set S if it is regular at every point $\mathbf{a}$ in S. f is **regular** if it is regular at $1_{\mathrm{loc}}(f)$. $\mathrm{Reg}(\overline{F(\Lambda)})$ denotes the set of regular rational functions.

Note that any rational function which is regular at a point $\mathbf{a}$ is also regular at the region containing $\mathbf{a}$. Another way of stating this condition, writing $f = \frac{h}{g}$ is to define a **leading Laurent monomial** of f to be of the form $\frac{h_i}{gj}$ where h_i is a dominant monomial of h and g_i is a dominant monomial of g. Of course f will have several leading Laurent monomials at $\mathbf{a}$ if $\mathbf{a}$ is a ghost root of h or g. The regularity condition is that f possesses some leading Laurent monomial $\neq \mathbf{1}_F$. For example, $f = \frac{\lambda_1+\lambda_2}{\lambda_1+\lambda_3}$ is regular at $(1,1,1)$ but not at $(2,1,1)$.

Lemma 7.4.2. *If f is regular, then any other generator f' of $\langle f\rangle$ is regular.*

Proof. If f' were not regular at $\mathbf{a}$, then its sole leading Laurent monomial at $\mathbf{a}$ would be $\mathbf{1}_F$, implying that some ν-neighborhood of $\mathbf{a}$ is in $1_{\mathrm{loc}}(f')$ and thus of $1_{\mathrm{loc}}(f)$, contrary to f possessing some leading Laurent monomial $\neq \mathbf{1}_F$. □

Lemma 7.4.3. *If $f, g \in \mathrm{Reg}(\overline{F(\Lambda)})$ such that $f \neq \mathbf{1}_F$ and $g \neq \mathbf{1}_F$, then the following elements are also in $\mathrm{Reg}(\overline{F(\Lambda)})$:*

$$f^*, \quad f^k \text{ with } k \in \mathbb{Z},\ f+g,\ |f|,\ f \wedge g.$$

Proof. Follows at once from Lemma 7.4.2 and Proposition 4.1.20. A more direct argument for $|f|$ is that writing $f = \frac{h}{g}$, we have

$$|f| = \frac{h^2+g^2}{gh},$$

but the leading Laurent monomial is either that of $h^2(gh)^* = f$ or that of $g^2(gh)^* = f^*$. □

Definition 7.4.4. The relation $f(\mathbf{a}) \cong_\nu \mathbf{1}_F$ is **regular** at $\mathbf{a}$ if f is regular at $\mathbf{a}$. Otherwise $f(\mathbf{a}) \cong_\nu \mathbf{1}_F$ is an **order relation**.

A 1^ν-set is **regular** if it can be written as $1_{\mathrm{loc}}(S)$ where each $f \in S$ is regular.

Lemma 7.4.5. *If $A = 1_{\mathrm{loc}}(f_1, ..., f_s)$ and $B = 1_{\mathrm{loc}}(g_1, ..., g_t)$ are regular then $A \cap B$ and $A \cup B$ are regular.*

Proof. Immediate from Lemma 5.1.6. □

Considering only dominant monomials at the neighborhood of $\mathbf{a}$, our local relations fall into two distinct cases:

- An order relation of the form $\mathbf{1}_F + g \cong_\nu \mathbf{1}_F$ with $g \in \overline{F(\Lambda)}$, i.e., $g(\mathbf{a}) \leq_\nu \mathbf{1}_F$. The resulting quotient ν-semifield† $\overline{F(\Lambda)}/\langle 1+g\rangle$ does not reduce the dimension at such a point $\mathbf{a}$, but only imposes new order relations on the variables.
- A regular relation; this reduces the dimensionality of the image of $\overline{F(\Lambda)}$ in the quotient ν-semifield†.

We aim to characterize those elements of $\overline{F(\Lambda)}$ that do not translate (locally) to order relations but only to regular relations (locally). This will allow us to characterize those relations which correspond to corner loci (tropical varieties in tropical geometry). The ν-kernels corresponding to these relations will be shown to form a sublattice of the lattice of principal ν-kernels (which is itself a sublattice of the lattice of ν-kernels).

Example 7.4.6. If $f \in \overline{F(\Lambda)}$ such that $f \neq f+1$ (i.e., 1 is essential in $f+1$), then $f+1 = \frac{f+1}{1}$ is not regular since 1 being essential in the numerator coincides with the denominator over some nonempty region.

Example 7.4.7. Consider the map $\phi : F(\lambda) \to F(\lambda)/\langle \lambda+1\rangle$ given by $\lambda \mapsto 1 \cong_\nu 1$. This map imposes the relation $\lambda+1 \cong_\nu 1$ on $F(\lambda)$, which is just the order relation $\lambda \leq_\nu 1$. Under the map ϕ, λ is sent to $\bar{\lambda} = \lambda\langle\lambda+1\rangle$, where now, in $\mathrm{Im}\,(\phi) = F(\bar{\lambda})$, $\bar{\lambda}$ and $\bar{1}$ are comparable, as opposed to the situation in $F(\lambda)$, where λ and 1 are not comparable. If instead of $\lambda+1$, we consider $|\lambda|+1 \cong_\nu \lambda+\lambda^{-1}+1$, then as $|\lambda| \geq_\nu 1$ the relation $|\lambda|+1 \cong_\nu 1$ means $|\lambda| \cong_\nu 1$, as yielded by the substitution map sending λ to 1. Note that $|\lambda|$ and 1 are comparable in $F(\lambda)$, since, as mentioned above, $|\lambda| \geq 1$, which is equivalent to the relation imposed by the equality $|\lambda|+1 \cong_\nu |\lambda|$ or equivalently by $|\lambda|^{-1}+1 \cong_\nu 1$. The ν-kernel $\langle|\lambda|^{-1}+1\rangle$, since $|\lambda|^{-1}+1 \cong_\nu 1$, is just the trivial ν-kernel $\langle 1\rangle = c$.

Note 7.4.8. As seen in Proposition 4.1.18 and Corollary 4.1.21, order relations affect the structures of ν-kernels in a ν-semifield†. For instance, consider the principal ν-kernel in a ν-semifield† F generated by an element $a \in F$. Then $b \notin \langle a\rangle$ for any element $b \in F$ such that b not comparable to a.

8. The role of the ν-kernel $\langle F\rangle$ of $\overline{F(\Lambda)}$ in the lattice of regular principal ν-kernels

Our overall objective in this section and the next is to describe the theory in terms of principal ν-kernels.

Definition 8.0.9. Denote the sublattice of principal ν-kernels of a ν-semifield† $\mathcal{S}$ by $\mathcal{P}(\mathcal{S})$.

Definition 8.0.10. A principal ν-kernel $K = \langle f\rangle$ of $\overline{F(\Lambda)}$ is **regular** if $f \in \overline{F(\Lambda)}$ is regular. In this case the 1-Set $1_{\mathrm{loc}}(f)$ corresponding to K also is called **regular**.

Proposition 8.0.11. *The set of regular principal ν-kernels forms a sublattice of $\mathcal{P}(\overline{F(\Lambda)})$.*

Proof. This follows directly from Corollary 4.1.30 and Lemma 7.4.3. □

8.0.1. *Corner loci and principal 1^ν-sets.*

Suppose $X \subseteq F^{(n)}$, and consider rational functions restricted to X. By Remark 5.1.5 and Lemma 5.1.6, the collection of corner loci is closed under intersections, while the collection of finitely generated principal corner loci is also closed under finite unions. Also, $\mathrm{C}_{\mathrm{loc}}(\emptyset) = X$ and $\mathrm{C}_{\mathrm{loc}}(\alpha) = \emptyset$.

By Lemma 7.4.5, the regular finitely generated principal corner loci comprise a sublattice, which we want to investigate.

Remark 8.0.12. By Remark 6.1.8 and Proposition 8.0.11 applied to the correspondence of Theorem 7.1.7, the lattice generated by principal corner internal ν-kernels with respect to (finite) products and intersections corresponds to the lattice of finitely generated corner loci. Note that $1_{\mathrm{loc}}(\mathbf{1}) = X$ and $1_{\mathrm{loc}}(\alpha) = \emptyset$ for $\alpha \not\cong_\nu \mathbf{1}$.

To delve deeper, we need to turn to the ν-kernel $\langle F\rangle$ of $\overline{F(\Lambda)}$.

Theorem 8.0.13. [28, Theorem (13.5.2)] *The lattice $\mathcal{P}(\langle F\rangle)$ is generated by the principal corner internal ν-kernels, and the sublattice of regular principal ν-kernels is generated by the regular, principal corner internal ν-kernels.*

The proof is rather long, requiring the concept of bounded ν-kernels, so we defer it until §9.

Corollary 8.0.14. *The lattice of (tangible) finitely generated corner loci corresponds to the lattice of principal (tangible) ν-kernels of $\langle F\rangle$. Intersections of (supertropical) hypersurfaces correspond to principal 1^ν-sets and ν-kernels, whereas intersections of tangible hypersurfaces correspond to tangible principal 1^ν-sets and ν-kernels.*

Proof. We use the correspondence of Remark 8.0.12, between principal (tangible) corner-loci and principal (tangible) corner internal ν-kernels of $\langle F\rangle$. □

Thus, supertropical varieties correspond to principal 1^ν-sets and ν-kernels, while tropical varieties correspond to regular principal 1^ν-sets and ν-kernels.

In order to bypass the ambiguity between ν-kernels and 1^ν-sets encountered in Example 6.2.1(i), we introduce one particular sub-ν-semifield† of $\overline{F(\Lambda)}$ of considerable interest.

8.1. The ν-kernel $\langle F\rangle$.

Assume that $F := (F, \mathcal{T}, \nu, \mathcal{G})$ is an archimedean ν-bipotent ν-semifield†.

Remark 8.1.1. $\langle\alpha\rangle = \langle\beta\rangle$ for any $\alpha, \beta \neq 1$ in F, in view of Proposition 4.1.18.

Definition 8.1.2. $\langle F\rangle$ denotes the ν-kernel given in Remark 8.1.1.

The ν-kernel $\langle F\rangle$ is preserved under any homomorphism ϕ for which $\phi(F) \neq 1$. In this subsection we show that $\langle F\rangle$ retains all the information in $\overline{F(\Lambda)}$ needed for the important family of principal 1^ν-sets and provides a 1:1 correspondence between ν-kernels and 1^ν-sets.

The ν-kernel $\langle F\rangle$ is much more sophisticated than what one might think at first blush, because of the convexity condition. Namely, any function lying between two constants is in the ν-kernel. For example, $3 \wedge |\lambda|$ is constant except in the interval between $\frac{1}{3}$ and 3, where it descends to 1 and then increases back to 3.

We begin by introducing the motivating example for this section.

Example 8.1.3. The 1^ν-set corresponding to the principal ν-kernel $\langle\lambda\rangle$ of $F(\lambda)$ is the set of all $\mathbf{a} \cong_\nu 1$. For any $\alpha \neq 1$ in F, we also have the principal ν-kernel $\langle\lambda\rangle \cap \langle\alpha\rangle = \langle|\lambda| \wedge |\alpha|\rangle = \left\langle \frac{|\lambda||\alpha|}{|\lambda|+|\alpha|} \right\rangle$. As $\lambda \notin \langle\lambda\rangle \cap \langle\alpha\rangle$, we conclude that $\langle\lambda\rangle \supset \langle\lambda\rangle \cap \langle\alpha\rangle$.

But $(|\lambda| \wedge |\alpha|)(\mathbf{a}) \cong_\nu 1$ iff $\mathbf{a} \cong_\nu 1$. It follows at once that $1_{\mathrm{loc}}(|\lambda| \wedge |\alpha|) = 1_{\mathrm{loc}}(\lambda)$.

In other words, different ν-kernels may have the same $1_{\mathrm{loc}}($. This ambiguity can be bypassed by intersecting all ν-kernels with $\langle F\rangle$.

Proposition 8.1.4. *If K is a $\mathcal{K}$-kernel of $\overline{F(\Lambda)}$, then $K' := K \cap \langle F\rangle$ is a $\mathcal{K}$-kernel of $\langle F\rangle$ satisfying $1_{\mathrm{loc}}(K) = 1_{\mathrm{loc}}(K')$.*

Proof. $K = \mathcal{K}ern(1_{\mathrm{loc}}(K))$ since K is a $\mathcal{K}$-kernel. Fix $\alpha \in F \setminus \{1\}$. By Proposition 4.1.28(ii) we have that $K \cap \langle F\rangle = K \cap \langle\alpha\rangle = \langle\{|f| \wedge |\alpha| : f \in K\}\rangle$. Now, for any $f \in K$, $f(\mathbf{a}) \cong_\nu 1$ for some $\mathbf{a} \in F^{(n)}$ if and only if $f(\mathbf{a}) \wedge |\alpha| = 1$ (since $|\alpha| > 1$) so $1_{\mathrm{loc}}(K') = 1_{\mathrm{loc}}(K)$. Thus

$$K' = K \cap \langle F\rangle = \mathcal{K}ern(1_{\mathrm{loc}}(K')) \cap \langle F\rangle = \mathcal{K}ern_{\langle F\rangle}(1_{\mathrm{loc}}(K')),$$

and so K' is a $\mathcal{K}$-kernel of $\langle F\rangle$. □

We want K' to be unique with this property, but for this we need to assume that F is complete and ν-archimedean, cf. Proposition 8.4.2 below.

Thus, we can pass down to $\langle F\rangle$, leading us to the next definition.

Definition 8.1.5. Define the equivalence relation

$$f \sim_{\langle F\rangle} f' \Leftrightarrow \langle f\rangle \cap \langle F\rangle = \langle f'\rangle \cap \langle F\rangle. \tag{8.1}$$

as ν-kernels of $\overline{F(\Lambda)}$. The equivalence classes are

$$[f] = \{f' \ : \ f' \text{ is a generator of } \langle f\rangle \cap \langle F\rangle\}.$$

Proposition 8.1.6. [28, Corollary (5.2.7)] *For any ν-kernel K of $\overline{F(\Lambda)}$, if $K \cap \langle F\rangle = \{1\}$, then $K = \{1\}$.*

Proof. For any $\mathbf{a} \in F^{(n)}$, $|f(\mathbf{a})| \wedge |\alpha| \cong_\nu 1$ if and only if $|f(\mathbf{a})| \cong_\nu 1$. □

Lemma 8.1.7.

$$\mathcal{P}(\langle F\rangle) = \{\langle f\rangle \cap \langle F\rangle : \langle f\rangle \in \mathcal{P}(\overline{F(\Lambda)})\}.$$

Proof. $\langle f\rangle \cap \langle F\rangle = \langle |f| \wedge |\alpha|\rangle \in \mathcal{P}(\langle F\rangle)$ for any $\langle f\rangle \in \mathcal{P}(\overline{F(\Lambda)})$. □

8.1.1. *Bounded rational functions.*

F is assumed to be an archimedean supertropical ν-semifield†.

Definition 8.1.8. A rational function $f \in \overline{F(\Lambda)}$ is said to be **bounded from below** if there exists some $\alpha >_\nu 1$ in F such that $|f| \geq_\nu \alpha$. f is said to be **bounded from above** (or **bounded**, for short) if there exists some $\alpha \geq_\nu 1$ in F such that $|f| \leq_\nu \alpha$.

Remark 8.1.9. [28, Remarks (5.1.3-4), (5.1.10-12)] Let $\langle f\rangle$ be a principal ν-kernel of $\overline{F(\Lambda)}$. Then

(i) f is bounded from below if and only if $\langle f\rangle \supseteq \langle\alpha\rangle$ for some $\alpha > 1$ in F. Moreover, any generator $g \in \langle f\rangle$ is bounded from below.
(ii) f is bounded if and only if $\langle f\rangle \subseteq \langle\alpha\rangle$ for some $\alpha \in F$. Moreover, any generator $g \in \langle f\rangle$ is bounded .

Definition 8.1.10. A principal ν-kernel $\langle f\rangle$ of $\overline{F(\Lambda)}$ is said to be **bounded from below**, if it is generated by a rational function bounded from below. $\langle f\rangle$ is said to be **bounded** if it is generated by a bounded rational function.

Proposition 8.1.11. [28, Corollary (5.1.8)] $1_{\mathrm{loc}}(\langle f\rangle) = \emptyset$ *if and only if f is bounded from below.*

Lemma 8.1.12. $\langle F\rangle = \{f \in \overline{F(\Lambda)} \ : \ f \text{ is bounded}\}$.

Proof. The assertion follows from Remark 8.1.9(ii), since $f \in \langle F\rangle$ if and only if $\langle f\rangle \subseteq \langle FW\rangle$. □

Lemma 8.1.13. *If F is complete, then $Fun(X, F)$ is complete, for any subset X of $F^{(n)}$.*

Proof. Suppose $\mathcal{S} \subset Fun(X, F)$ is bounded from below, say by $h \in Fun(X, F)$. Then for any $\mathbf{a} \in X$ the set $W := \{f(\mathbf{a}) : f \in X\}$ is bounded from below by $h(\mathbf{a})$ and thus has an infimum $\bigwedge_{f\in W} f(\mathbf{a})$. It is readily seen that the function $g \in Fun(F^{(n)}, F)$ defined by $g(\mathbf{a}) = \bigwedge_{f\in W} f(\mathbf{a})$ is an infimum for X, i.e., $g = \bigwedge_{f\in W} f$. Analogously, if W is bounded then $(\bigvee_{f\in W} f)(\mathbf{a}) = \bigvee_{f\in W} f(\mathbf{a})$ is the supremum of W. □

8.1.2. *The map ω.*

Definition 8.1.14. Fixing $\alpha >_\nu 1$, define the map $\omega : \overline{F(\Lambda)}^+ \to \langle F\rangle^+$ by

$$\omega(|f|) = |f| \wedge |\alpha|.$$

Lemma 8.1.15. *ω is a lattice homomorphism.*

Proof. Indeed, $\omega(|f| + |g|) = (|f| + |g|) \wedge |\alpha| = (|f| \wedge |\alpha|) + (|g| \wedge |\alpha|) = \omega(|f|) + \omega(|g|)$ and $\omega(|f| \wedge |g|) = (|f| \wedge |g|) \wedge |\alpha| = (|f| \wedge |\alpha|) \wedge (|g| \wedge |\alpha|) = \omega(|f|) \wedge \omega(|g|)$.

□

ω induces a lattice map

$$\Omega : \mathcal{P}(\overline{F(\Lambda)}) \to \mathcal{P}(\overline{F(\Lambda)}) \cap \langle F\rangle$$

such that $\Omega(\langle f\rangle) = \langle \omega(|f|)\rangle = \langle |f| \wedge |\alpha|\rangle = \langle f\rangle \cap \langle F\rangle$.

Lemma 8.1.16. (i) *Ω preserves both intersections and products of ν-kernels.*

(ii) *If $\langle f\rangle$ is a ν-kernel that is bounded from below, then $\langle f\rangle \cap \langle F\rangle = \langle F\rangle$. In fact, every principal ν-kernel whose 1^ν-set is the empty set is mapped to $\langle F\rangle$.*

(iii) $1_{\mathrm{loc}}(\Omega(\langle f\rangle)) = 1_{\mathrm{loc}}(\langle f\rangle \cap \langle H\rangle) = 1_{\mathrm{loc}}(\langle f\rangle) \cup \emptyset = 1_{\mathrm{loc}}(\langle f\rangle)$.

(iv) *Any ν-subkernel $K \subset \langle \alpha\rangle$ must satisfy $K \cap F = \{1\}$.*

Proof. (i) By Proposition 5.2.7.

(ii) By Remark 8.1.9 and Lemma 8.1.7.

(iii) $1_{\mathrm{loc}}(\langle \alpha\rangle) = \emptyset$ and $1_{\mathrm{loc}}(\langle f\rangle \cap \langle g\rangle) = 1_{\mathrm{loc}}(\langle f\rangle) \cup 1_{\mathrm{loc}}(\langle g\rangle)$, for any principal ν-kernel $\langle f\rangle$.

(iv) Any $\alpha \neq 1$ generates $\langle F\rangle$. □

We now show that restricting 1_{loc} and $\mathcal{K}ern$ to $\langle F\rangle$ does not affect the 1^ν-sets and that each $\mathcal{K}$-kernel of $\overline{F(\Lambda)}$ restricts to a $\mathcal{K}$-kernel in $\langle F\rangle$.

In Proposition 4.1.27, we have shown using the substitution homomorphism ψ that any point $a = (\alpha_1, ..., \alpha_n) \in F^{(n)}$ corresponds to the maximal ν-kernel

$$\left\langle \frac{\lambda_1}{\alpha_1}, ..., \frac{\lambda_n}{\alpha_n} \right\rangle = \left\langle \left|\frac{\lambda_1}{\alpha_1}\right| + ... + \left|\frac{\lambda_n}{\alpha_n}\right| \right\rangle = \left\langle \left|\frac{\lambda_1}{\alpha_1}\right| \cdots \left|\frac{\lambda_n}{\alpha_n}\right| \right\rangle = \left\langle \frac{\lambda_1}{\alpha_1} \right\rangle \cdots \left\langle \frac{\lambda_n}{\alpha_n} \right\rangle.$$

For example, consider the homomorphism $\psi : F(\lambda_1) \to F$ given by $\lambda_1 \mapsto 1$. By Theorem 4.1.10, the ν-kernel of its restriction $\psi|_{\langle F\rangle} : \langle F\rangle \to \psi(\langle F\rangle) = F$ is $\mathcal{K}ern(\psi|_{\langle F\rangle}) = \mathcal{K}ern(\psi) \cap \langle F\rangle = \langle \lambda_1\rangle \cap \langle F\rangle$. Thus, the proposition applies to $\langle F\rangle$ providing the maximal ν-kernel $\langle x\rangle \cap \langle F\rangle$. We will now show that any maximal ν-kernel of $\langle F\rangle$ has that form.

Proposition 8.1.17. *If K is a maximal ν-kernel of $\langle F\rangle$, then*

$$K = \Omega\left(\left\langle \frac{\lambda_1}{\alpha_1}, ..., \frac{\lambda_n}{\alpha_n} \right\rangle\right)$$

for suitable $\alpha_1, ..., \alpha_n \in F$.

Proof. Denote $L_a = (|\frac{\lambda_1}{\alpha_1}| + ... + |\frac{\lambda_n}{\alpha_n}|) \wedge |\alpha|$ with $\alpha \neq 1$, for $a = (\alpha_1, ..., \alpha_n)$. By Lemma 8.1.16 we may assume that $1_{\mathrm{loc}}(K) \neq \emptyset$, since the only ν-kernel corresponding to the empty set is $\langle F\rangle$ itself. If $a \in 1_{\mathrm{loc}}(K)$, then as $1_{\mathrm{loc}}(L_a) = \{a\} \subseteq 1_{\mathrm{loc}}(K)$, we have that $\langle L_a\rangle \supseteq K$. Thus, the maximality of K implies that $K = \langle L_a\rangle$. □

8.2. **Completions of idempotent ν-semifields†.**

We recall Definition 2.2.8 of "complete," viewing a ν-semifield† as a lattice. In essence it is enough to consider $\mathscr{R}$ from Example 3.1.6.

8.2.1. *Kernels of ν-archimedean idempotent ν-semifields†.*

Theorem 8.2.1. [36, Theorem 2.3.10 (Hölder)] *The following statements are equivalent for an idempotent ν-semifield† $\mathcal{S}$.*

(1) *$\mathcal{S}$ is simple.*

(2) *$\mathcal{S}$ is totally ordered and archimedean.*

(3) *$\mathcal{S}$ can be embedded into the max-plus algebra of $(\mathbb{R}, +)$.*

The supertropical version: Any ν-archimedean supertropical ν-semifield† $\mathcal{S} = (\mathcal{S}, \mathcal{T}, \nu, \mathcal{G})$ can be embedded into $\mathscr{R}$.

Corollary 8.2.2. [36] *Any complete divisible totally ordered ν-archimedean supertropical ν-semifield† is isomorphic to $\mathscr{R}$.*

To avoid duplication for ν-kernels having the same 1^ν-set, our next step is to work with $\mathscr{R}$ and restrict to ν-subkernels of $\langle\mathscr{R}\rangle$ in $\mathscr{R}(\Lambda)$.

Definition 8.2.3. A subset A of the poset P is called **co-initial** in P if for every $x \in P$ there exists some $f \in A$ such that $f \leq x$. A is **co-final** in P if for every $x \in P$ there exists some $f \in A$ such that $x \leq f$.

Definition 8.2.4. The subset A of the idempotent ν-semifield† $\mathcal{S}$ is called **left dense** in $\mathcal{S}$ if A^+ (cf. Definition 3.3.2) is co-initial in $\mathcal{S}^+$, and A is called **right dense** in H if A^+ is co-final in $\mathcal{S}^+$.

Definition 8.2.5. A **completion** of the idempotent ν-semifield† $\mathcal{S}$ is a pair (H, θ) where H is a complete idempotent ν-semifield† and $\theta : \mathcal{S} \to H$ is a monomorphism whose image is dense (left and right) in H.

Remark 8.2.6. Each ν-archimedean idempotent ν-semifield† has a unique completion up to isomorphism. The proof follows the standard lines, with details given in [36, Theorem (2.3.4)].

Remark 8.2.7. When the supertropical ν-semifield† F is ν-archimedean, any ν-kernel K of $\overline{F(\Lambda)}$ also has a completion, which we denote as $\bar{K} \subseteq \overline{F(\Lambda)}$, which by Corollary 8.2.9 is a ν-kernel of $\overline{F(\Lambda)}$.

Theorem 8.2.8. [36, Theorem 2.3.6] *Let $\mathcal{S}$ be an ν-archimedean sub-ν-semifield† of a complete idempotent ν-semifield† $\hat{\mathcal{S}}$. Then the following are equivalent:*

1. *$\hat{\mathcal{S}}$ is the completion of $\mathcal{S}$.*
2. *$\mathcal{S}$ is left dense in $\hat{\mathcal{S}}$, and if A is an idempotent subsemifield of $\hat{\mathcal{S}}$ that is complete and contains $\mathcal{S}$, then $A = \hat{\mathcal{S}}$.*

Corollary 8.2.9. [36] *Suppose $\mathcal{S}$ is a left dense ν-archimedean idempotent sub-ν-semifield† of the complete idempotent ν-semifield† $\hat{\mathcal{S}}$. Then the ν-kernel of $\hat{\mathcal{S}}$ generated by a ν-kernel K of $\mathcal{S}$, is the completion of K in $\hat{\mathcal{S}}$.*

By Remark 8.2.6, $\mathscr{R}(\Lambda)$ has a unique completion to a complete ν-archimedean idempotent ν-semifield† $\overline{\mathscr{R}(\Lambda)}$ in $Fun(\mathscr{R}^{(n)}, \mathscr{R})$. By Theorem 8.2.8, $\mathscr{R}(\Lambda)$ is dense in $\overline{\mathscr{R}(\Lambda)}$.

8.3. The ν-kernel $\langle\mathscr{R}\rangle$ of $\mathscr{R}(\Lambda)$ as bounded functions.

Proposition 8.3.1. [28, Proposition (5.2.1)] *For any principal bounded ν-kernel $\langle f\rangle \in \mathcal{P}(\langle\mathscr{R}\rangle)$, there exists an unbounded ν-kernel $\langle f'\rangle \in \mathcal{P}(\mathscr{R}(\Lambda))$, such that*

$$\langle f\rangle = \langle f'\rangle \cap \langle\mathscr{R}\rangle.$$

In particular, $\langle f'\rangle \supset \langle f\rangle$ and $1_{\mathrm{loc}}(f') = 1_{\mathrm{loc}}(f)$.

Proof. Suppose $f \in \langle\mathscr{R}\rangle$ is bounded. Then there exists some $\beta_1 \in \mathscr{R}$ such that $|f(\mathbf{a})| \leq_\nu \alpha_1 \in \mathscr{R}$. Similarly for each $2 \leq i \leq n$ there exists some $\beta_i \in \mathscr{R}$ such that $|f(\mathbf{a})| \leq_\nu \alpha_i$ whenever $\alpha_i \geq_\nu \beta_i$. As $|f|$ is continuous we may assume that $\alpha_i = \alpha$ are all the same. Now define the function

$$f' = |(\beta^{-1}|\lambda_1| \wedge .... \wedge \beta^{-1}|\lambda_n| + 1)| + |f(\Lambda)|$$

where $\beta = \sum_{i=1}^n |\beta_i|$. Write $g(\lambda_1, ...\lambda_n) = \beta^{-1}|\lambda_1| + ... + \beta^{-1}|\lambda_n| + 1$. Let

$$S = \{\mathbf{a} = (a_1, \ldots, a_n) \in F^{(n)} : |a_i| > \beta \ \forall i\}.$$

Then for every $a \in S$, $f'(a) = g(a) + |\alpha|$. Moreover, for every $b = (b_1, ..., b_n) \notin S$ there exists some j such that $|b_j| \leq \beta$ thus we have that $(\beta^{-1}|b_1| + ... + \beta^{-1}|b_n|) \leq 1$ and so $g(b) = 1$. By construction $1_{\mathrm{loc}}(f) \subseteq 1_{\mathrm{loc}}(g)$, so $1_{\mathrm{loc}}(f') = 1_{\mathrm{loc}}(|g| + |f|) = 1_{\mathrm{loc}}(g) \cap 1_{\mathrm{loc}}(f) = 1_{\mathrm{loc}}(f)$. Finally f' is not bounded since $|g|$ is not bounded and $|f'| = |g| + |f| \geq_\nu |g|$. Now, as $|f'| = |g| + |f|$ we have that $|f| \leq |f'|$, so $f \in \langle f'\rangle$. On the other hand, since f' is not bounded, clearly $f' \notin \langle f\rangle$. Finally, $g(a) \geq_\nu 1$ for any $a \in S$. Thus, $f'(a) \wedge |\alpha| \cong_\nu (g(a) + |\alpha|) \wedge |\alpha| \cong_\nu |\alpha|$, while for $a \notin S$ $f'(a) \wedge |\alpha| \cong_\nu (g(a) + |f(a)|) \wedge |\alpha| \cong_\nu (1 + |f(a)|) \wedge |\alpha| \cong_\nu |f(a)| \wedge |\alpha| \cong_\nu |f(a)|$, since $|f| \leq_\nu |\alpha|$. So we get that $|f'| \wedge |\alpha| \cong_\nu |f|$ which means that $\langle f\rangle = \langle f'\rangle \cap \langle\mathscr{R}\rangle$. (Note that $f' = |f'|$ by definition, since $f' \geq_\nu 1$.) □

8.4. **Principal 1^ν-sets and bounded ν-kernels.**

It turns out that $\langle\mathscr{R}\rangle$ possesses enough distinct bounded copies of the principal ν-kernels of $\mathscr{R}(\Lambda)$ to represent faithfully the principal 1^ν-sets.

Remark 8.4.1. The restriction of the image of the operator $\mathcal{K}ern : \mathbf{2}^{\mathscr{R}^{(n)}} \to \mathbf{2}^{\mathscr{R}(\Lambda)}$ to $\mathbf{2}^{\langle\mathscr{R}\rangle}$ is

$$\mathcal{K}ern_{\langle\mathscr{R}\rangle}(Z) = \{f \in \langle\mathscr{R}\rangle \ : \ f(a_1,...,a_n) = 1, \ \forall (a_1,...,a_n) \in Z\} = \mathcal{K}ern(Z) \cap \langle\mathscr{R}\rangle. \tag{8.2}$$

Furthermore, the assertions of this subsection apply to $\mathcal{K}ern_{\langle\mathscr{R}\rangle}$ and the restriction $1_{\mathrm{loc}}|_{\langle\mathscr{R}\rangle} : \mathbf{2}^{\langle\mathscr{R}\rangle} \to \mathbf{2}^{\mathscr{R}^{(n)}}$ of 1_{loc} to $\mathbf{2}^{\langle\mathscr{R}\rangle}$.

When there is no ambiguity, we denote $\mathcal{K}ern_{\langle\mathscr{R}\rangle}$ and $1_{\mathrm{loc}}|_{\mathbf{2}^{\langle\mathscr{R}\rangle}}$ respectively as $\mathcal{K}ern$ and 1_{loc}.

Applying Lemma 8.1.16 to the ν-semifield† $\mathscr{R}$, we see that

$$\Omega : \mathcal{P}(\mathscr{R}(\Lambda)) \to \mathcal{P}(\langle\mathscr{R}\rangle)$$

is a lattice homomorphism of the lattice $(\mathcal{P}(\mathscr{R}(\Lambda)), \cdot, \cap)$ onto $(\mathcal{P}(\langle\mathscr{R}\rangle), \cdot, \cap)$, such that $1_{\mathrm{loc}}(\langle f\rangle) = 1_{\mathrm{loc}}(\Omega(\langle f\rangle))$.

Let $f \in \mathcal{P}(\langle\mathscr{R}\rangle)$ and let $A = \{g \in \mathscr{R}(\Lambda) : \Omega(\langle g\rangle) = f\}$. Define the ν-kernel $K = \langle A\rangle$ of $\mathscr{R}(\Lambda)$. Then by Remark 8.1.16, if $g \in K$ then $1_{\mathrm{loc}}(g) = 1_{\mathrm{loc}}(f)$.

We aim for the 1:1 correspondence

$$\langle f\rangle \in \mathcal{P}(\langle\mathscr{R}\rangle) \leftrightarrow 1_{\mathrm{loc}}(f)$$

between the principal 1^ν-sets in $\mathscr{R}^{(n)}$ and the ν-kernels in $\mathcal{P}(\langle\mathscr{R}\rangle)$.
If $1_{\mathrm{loc}}(g) = 1_{\mathrm{loc}}(f)$, then $1_{\mathrm{loc}}(\Omega(g)) = 1_{\mathrm{loc}}(f)$ since $1_{\mathrm{loc}}(g) = 1_{\mathrm{loc}}(\Omega(g))$, and $\Omega(g), f \in \mathcal{P}(\langle\mathscr{R}\rangle)$. Thus in view of the above $\Omega(\langle g\rangle) = \langle f\rangle$. Consequently $1_{\mathrm{loc}}(K) = 1_{\mathrm{loc}}(f)$, and K is the maximal ν-kernel of $\mathscr{R}(\Lambda)$ having this property. Our next result justifies the use of $\langle\mathscr{R}\rangle$.

Proposition 8.4.2. *If $\langle f\rangle \subseteq \langle\mathscr{R}\rangle$ and $h \in \langle f\rangle$ is such that $1_{\mathrm{loc}}(h) = 1_{\mathrm{loc}}(f)$, then h is a generator of $\langle f\rangle$.*

Proof. The assertion is obvious in the special case for which $\langle f\rangle \cong_\nu \{1\}$. So, as $1_{\mathrm{loc}}(h) = 1_{\mathrm{loc}}(f)$ we may assume that f and h are not ν-equivalent to 1. If $\langle f\rangle = \langle\alpha\rangle$, then $1_{\mathrm{loc}}(h) = 1_{\mathrm{loc}}(f) = \emptyset$ implies by Lemma 8.1.16 (3) that $\langle h\rangle = \langle\mathscr{R}\rangle = \langle f\rangle$.

Suppose there is a rational function $h \in \langle f\rangle$ which is not a generator of $\langle f\rangle$ but satisfying $1_{\mathrm{loc}}(h) = 1_{\mathrm{loc}}(f)$. By Corollary 4.1.21, for each $k \in \mathbb{N}$ there exists some $\mathbf{a} \in \mathscr{R}^n$ for which $|f(\mathbf{a})| >_\nu |h(\mathbf{a})|^k$. For any $k \in \mathbb{N}$, define the set $U_k = \{\mathbf{a} \ : \ |f(\mathbf{a})| >_\nu |h(\mathbf{a})|^k\}$. As $\mathscr{R}$ is dense, for any $\mathbf{a} \in U_k$ there exists a neighborhood $B_{\mathbf{a}} \subset U_k$ containing $\mathbf{a}$ such that for all $\mathbf{a}' \in B_{\mathbf{a}}$, $|f(\mathbf{a}')| > |h(\mathbf{a}')|^k$. Now, since h and f are bounded, the U_k are bounded regions inside $\mathscr{R}^{(n)}$. Taking the closures, we may assume that the U_k are closed (nonempty). Since $1_{\mathrm{loc}}(h) = 1_{\mathrm{loc}}(f)$, $|f(\mathbf{a})| >_\nu |h(\mathbf{a})|^k$ implies that $|h(\mathbf{a})|, |f(\mathbf{a})| >_\nu 1$, so, by the definition of U_k we get the sequence of strict inclusions $U_1 \supset U_2 \supset \cdots \supset U_k \supset \ldots$. Thus, since $\mathscr{R}$ is complete, there exists an element $\mathbf{b} \in \mathrm{B} = \bigcap_{\mathbb{N}} B_k$. Now, for $\mathbf{a} \notin 1_{\mathrm{loc}}(h)$, $|h(\mathbf{a})| > 1$, and thus there exists some $r = r(\mathbf{a}) \in \mathbb{N}$ such that $|h(\mathbf{a})|^r >_\nu |f(\mathbf{a})|$ and thus $\mathbf{a} \notin \mathrm{B}$ thus $\mathbf{b} \notin \mathscr{R}^{(n)} \setminus 1_{\mathrm{loc}}(h)$. On the other hand, if $\mathbf{b} \in 1_{\mathrm{loc}}(h)$ then $\mathbf{b} \in 1_{\mathrm{loc}}(f)$ so $1 \cong_\nu |f(\mathbf{b})| \leq_\nu |h(\mathbf{b})| \cong_\nu 1$. Thus $\bigcap_{\mathbb{N}} B_k = \emptyset$, contradiction. □

Using Theorem 5.3.6 and Proposition 8.1.4, we conclude:

Theorem 8.4.3. *There is a $1:1$, order-reversing correspondence*

$$\{\textit{principal } 1^\nu\textit{-sets of } \mathscr{R}^n\} \leftrightarrow \{\mathcal{K} - \textit{principal } \nu\textit{-kernels of } \langle\mathscr{R}\rangle\}, \tag{8.3}$$

given by $Z \mapsto \mathcal{K}ern_{\langle\mathscr{R}\rangle}(Z)$; the reverse map is given by $K \mapsto 1_{\mathrm{loc}}(K)$.

Proof. Every principal ν-kernel gives rise to a principal 1^ν-set by the definition of 1_{loc}. The reverse direction follows from Proposition 8.4.2, as every principal ν-kernel which produces a principal 1^ν-set via 1_{loc} is in fact a $\mathcal{K}$-kernel. □

Proposition 8.4.4. *Let $\langle f\rangle$ be a principal ν-kernel in $\mathcal{P}(\langle\mathscr{R}\rangle)$. Then $\langle f\rangle$ is a $\mathcal{K}$-kernel.*

Proof. We need to show that $\mathcal{K}ern(1_{\mathrm{loc}}(f)) \subseteq \langle f\rangle$.
Let $h \in \langle\mathscr{R}\rangle$ such that $h \in \mathcal{K}ern(1_{\mathrm{loc}}(f))$. Then $h(x) = 1$ for every $x \in 1_{\mathrm{loc}}(f)$ and so $1_{\mathrm{loc}}(f) \subseteq 1_{\mathrm{loc}}(h)$. If $|h| \leq |f|^k$ for some $k \in \mathbb{N}$ then $h \in \langle f\rangle$. Thus in particular we may assume that $h \neq 1$. Now, by Corollary 5.2.12, $1_{\mathrm{loc}}(\langle f\rangle \cap \langle h\rangle) = 1_{\mathrm{loc}}(f) \cup 1_{\mathrm{loc}}(h) = 1_{\mathrm{loc}}(h)$. Since $h \neq 1$, $1_{\mathrm{loc}}(h) \neq \mathscr{R}^{(n)}$ and thus $\langle f\rangle \cap \langle h\rangle \neq \{1\}$. Again, $1_{\mathrm{loc}}(\langle f\rangle\langle h\rangle) = 1_{\mathrm{loc}}(f) \cap 1_{\mathrm{loc}}(h) = 1_{\mathrm{loc}}(f)$ by Corollary 5.2.12.

Thus $\langle f, h\rangle = \langle f\rangle\langle h\rangle \neq \langle\mathscr{R}\rangle$ for otherwise $1_{\mathrm{loc}}(f) = \emptyset$. Consequently for $g = |f| \wedge |h|$, the ν-kernel $K = \langle g\rangle = \langle f\rangle \cap \langle h\rangle$ satisfies $\{1\} \neq K \subseteq \langle f\rangle$, implying $g \in \langle f\rangle$ and $1_{\mathrm{loc}}(g) = 1_{\mathrm{loc}}(h)$. Thus by Proposition 8.4.2, g is a generator of $\langle h\rangle$, so $\langle h\rangle = K \subseteq \langle f\rangle$ as desired. □

8.5. **The wedge decomposition in $\langle F\rangle$.**

Lemma 8.5.1. *Suppose $f=\sum_{i=1}^k f_i\in F[x_1,...,x_n]$ is a supertropical polynomial written as the sum of its monomials, then $\tilde{f}:=\bigwedge_{i=1}^k \widehat{f}_i$ is corner internal.*

Proof. $\tilde{f}\geq 1$. For any given $i\in\{1,...,k\}$, then Remark 7.3.10 yields $1_{\rm loc}(\Phi_{CI}(|\widehat{f}_i|))=1_{\rm loc}(\Phi_{CI}(\widehat{f}_i))=1_{\rm loc}(\widehat{f}_i)$. Then

$$1_{\rm loc}(\Phi_{CI}(\tilde{f}))=1_{\rm loc}(\bigwedge_{i=1}^k \Phi_{CI}(|\widehat{f}_i|))=1_{\rm loc}(\bigwedge_{i=1}^k \widehat{f}_i)$$

so $\Phi_{CI}(\tilde{f})=\tilde{f}=\bigwedge_{i=1}^k \widehat{f}_i$, which by Theorem 7.3.9 is corner internal. □

Definition 8.5.2. A **wedge decomposition** of a rational function $f\in F(x_1,...,x_n)$ is an expression

$$|f|=\bigwedge_i |u_i| \tag{8.4}$$

for $u_i\in F(x_1,...,x_n)$. This wedge decomposition is **associated to** a ν-kernel intersection $\langle f\rangle=\bigcap K_i$ if each $K_i=\langle u_i\rangle$.

Note by Lemmas 5.2.2(ii) and 5.2.3 that $1_{\rm loc}(f)=\bigcup_i 1_{\rm loc}(u_i)$.

Theorem 8.5.3. *Any intersection $\langle f\rangle=\bigcap K_i$ of (principal) ν-subkernels of $\langle\mathscr{R}\rangle$ has an associated wedge decomposition.*

Proof. It is enough to prove the assertion when $\langle f\rangle=\langle u\rangle\cap\langle v\rangle$ for $u,v\in\langle\mathscr{R}\rangle$, and then to conclude by induction. Let $\bar{f}=|u|\wedge|v|$. $\langle\bar{f}\rangle=\langle f\rangle$, so there exist some $q_1,...,q_k\in\mathscr{R}(\Lambda)$ such that $\sum_{i=1}^k q_i=1$ and $|f|=\sum_{i=1}^k q_i\bar{f}^{d(i)}$ with $d(i)\in\mathbb{Z}_{\geq 0}$ ($d(i)\geq 0$ since $|f|\geq 1$). Thus

$$f=\sum_{i=1}^k q_i(|u|\wedge|v|)^{d(i)}=\left(\sum_{i=1}^k q_i|u|^{d(i)}\right)\bigwedge\left(\sum_{i=1}^k q_i|v|^{d(i)}\right)$$

$$=|g|\wedge|h|$$

where $g=|g|=\sum_{i=1}^k q_i|u|^{d(i)}$ and $h=|h|=\sum_{i=1}^k q_i|v|^{d(i)}$. Thus

$$1_{\rm loc}(f)\supseteq 1_{\rm loc}(g)\supseteq 1_{\rm loc}(u);\qquad 1_{\rm loc}(f)\supseteq 1_{\rm loc}(h)\supseteq 1_{\rm loc}(v).$$

We claim that $|g|$ and $|h|$ generate $\langle|u|\rangle$ and $\langle|v|\rangle$, respectively. Since $\langle\bar{f}\rangle=\langle f\rangle$, we see that $\bar{f}(\mathbf{a})=1\Leftrightarrow|f|(\mathbf{a})=1$ for any $\mathbf{a}\in\mathscr{R}^{(n)}$. Let $q_j\bar{f}^{d(j)}$ be a dominant term of $|f|$ at $\mathbf{a}$, i.e.,

$$|f(\mathbf{a})|\cong_\nu\sum_{i=1}^k q_i(\mathbf{a})(\bar{f}(\mathbf{a}))^{d(i)}\cong_\nu q_j(\mathbf{a})(\bar{f}(\mathbf{a}))^{d(j)}.$$

Then $\bar{f}(\mathbf{a})\cong_\nu 1\Leftrightarrow q_j(\mathbf{a})(\bar{f}(\mathbf{a}))^{d(j)}\cong_\nu 1$. Hence, $\bar{f}(\mathbf{a})\cong_\nu 1\Leftrightarrow q_j(\mathbf{a})=1$.

Now, consider $\mathbf{a}\in 1_{\rm loc}(g)$. Then $g(\mathbf{a})=\sum_{i=1}^k q_i|u|^{d(i)}\cong_\nu 1$. Let $q_t|u|^{d(t)}$ be a dominant term of g at $\mathbf{a}$. If $q_t(\mathbf{a})\cong_\nu 1$ then $|u|^{d(t)(\mathbf{a})}\cong_\nu 1$ and thus $\mathbf{a}\in 1_{\rm loc}(u)$.

So we may assume that $q_t(\mathbf{a})<_\nu 1$ (since $\sum_{i=1}^k q_i\cong_\nu 1$) and so, as above, $q_t\bar{f}^{d(t)}$ is not a dominant term of $|f|$ at $\mathbf{a}$. Hence,

$$|u(\mathbf{a})|^{d(j)}=q_j(\mathbf{a})|u(\mathbf{a})|^{d(j)}<_\nu q_t(\mathbf{a})|u(\mathbf{a})|^{d(t)}\cong_\nu g(\mathbf{a})\cong_\nu 1,$$

for any dominant term of $|f|$ at $\mathbf{a}$. Thus

$$q_j(\mathbf{a})(\bar{f}(\mathbf{a}))^{d(j)}=q_j(\mathbf{a})(|u|(\mathbf{a})\wedge|v|(\mathbf{a}))^{d(j)}\leq q_j(\mathbf{a})|u(\mathbf{a})|^{d(j)}<_\nu 1. \tag{8.5}$$

On the other hand, $\bar{f}(\mathbf{a})\cong_\nu 1$ since $1_{\rm loc}(f)\supseteq 1_{\rm loc}(g)$, and thus $q_j(\mathbf{a})(\bar{f}(\mathbf{a}))^{d(j)}\cong_\nu 1$, contradicting (11.3). Hence $1_{\rm loc}(g)\subseteq 1_{\rm loc}(u)$, so, by the above, $1_{\rm loc}(g)=1_{\rm loc}(u)$, which in turn yields that g is a generator of $\langle|u|\rangle=\langle u\rangle$, by Proposition 8.4.2. The proof for h and $|v|$ is analogous. Consequently, we have that $1_{\rm loc}(g)=1_{\rm loc}(|g|)=1_{\rm loc}(|u|)$ and $1_{\rm loc}(h)=1_{\rm loc}(|v|)$, as desired. □

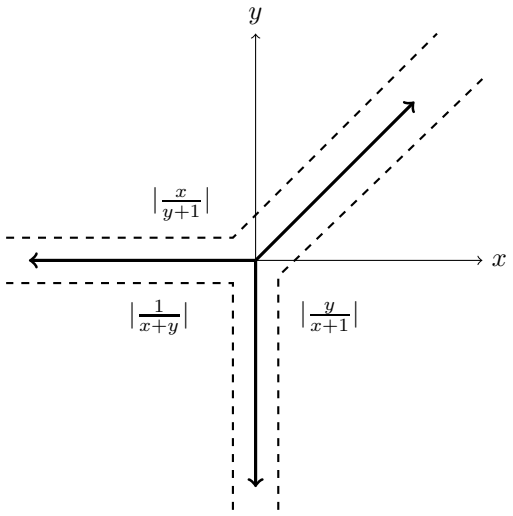

FIGURE 1. $\tilde{f} = |\frac{x}{y+1}| \wedge |\frac{y}{x+1}| \wedge |\frac{1}{x+y}|$

8.6. **Example: The tropical line revisited.**

For ease of notation, we write x for λ_1 and y for λ_2.

Note 8.6.1. In the following example we consider the rational function

$$\widehat{f} = \left| \frac{x}{y+1} + \frac{y}{x+1} + \frac{1}{x+y} \right| \wedge |\alpha| \in \langle \mathscr{R} \rangle$$

for any $\alpha \in \mathscr{R} \setminus \{1\}$.

Example 8.6.2. Let $f = x+y+1$ be the tropical line, considered already in Example 6.2.1(i). Its corresponding 1^ν-set is defined by the rational function $\widehat{f} = \frac{x}{y+1} + \frac{y}{x+1} + \frac{1}{x+y}$, and so, its corresponding ν-kernel in $\mathscr{R}(x,y)$ is $\langle \widehat{f} \rangle$ and the bounded copy in $\langle \mathscr{R} \rangle$ is $\langle |\widehat{f}| \wedge |\alpha| \rangle = \langle \widehat{f} \rangle \cap \langle \mathscr{R} \rangle$. As mentioned above $\widehat{f} \sim_{\langle \mathscr{R} \rangle} \left| \frac{x}{y+1} \right| \wedge \left| \frac{y}{x+1} \right| \wedge \left| \frac{1}{x+y} \right|$, and any of the three above terms can be omitted. Thus we have that

$$\langle \widehat{f} \rangle \cap \langle \mathscr{R} \rangle = \left\langle \frac{x}{y+1} \right\rangle \cap \left\langle \frac{y}{x+1} \right\rangle \cap \langle x+y \rangle \cap \langle \mathscr{R} \rangle,$$

where each of the ν-kernels comprising the intersection (excluding $\langle \mathscr{R} \rangle$) is contained in each of the other ν-kernels. (In the third ν-kernel from the left we choose to take $x+y$ as a generator instead of its inverse.) Now, taking logarithms, it can be seen that $1_{\mathrm{loc}}(\langle x+y \rangle)$ is precisely the union of the bounding rays of the third quadrant. Furthermore,

$$\langle x+y \rangle = \langle |x+y| + |x| \rangle \cap \langle |x+y| + |y| \rangle,$$

since

$$x + y \sim_{\langle \mathscr{R} \rangle} |x+y| = |x+y| + (|x| \wedge |y|) = (|x+y| + |x|) \wedge (|x+y| + |y|)\,.$$

This wedge decomposition of $|x+y|$ is quite natural. The geometric locus of the equation $|x| \wedge |y| = 1$ in logarithmic scaling is precisely the union of the x-axis corresponding to $|x| \cong_\nu 1$ and the y-axis corresponding to $|y| \cong_\nu 1$.

Using the wedge decomposition, we can define any segment and ray in $\mathscr{R}^2$ by means of principal ν-kernels, so the only irreducible 1^ν-sets turn out to be the points in the plane. This does not hamper us in developing geometry, since we may also restrict our attention to sublattices of the lattice of principal ν-kernels.

9. Hyperspace-ν-kernels and region ν-kernels

In order to prove Theorem 8.0.13, we separate principal ν-kernels into two classes of ν-kernels, HO-kernels and bounded from below ν-kernels.The latter ν-kernels are eliminated by passing to the ν-kernel $\langle F \rangle$. Using these classes of ν-kernels we introduce the hyperspace-region decomposition given in Theorem 9.4.8.

9.1. **Hyperspace-kernels.**

Remark 9.1.1. Though we consider the ν-semifield† of rational functions $F(\lambda_1, ..., \lambda_n)$, most of the results introduced in this section are applicable to any finitely generated sub-ν-semifield† of $\overline{F(\Lambda)}$ over F. In particular, $\langle F\rangle \subset F(\lambda_1, ..., \lambda_n)$ is just another case of a finitely generated ν-semifield† over F, taking the generators $\lambda_i \wedge |\alpha|$ for $1 \leq i \leq n$, and $\alpha \in F \setminus \{1\}$. In this case, the generators are bounded, and we specifically designate the results that are true only for unbounded generators.

Definition 9.1.2. An **$\mathscr{L}$-monomial** is a non-constant Laurent monomial $f \in F(\lambda_1, ..., \lambda_n)$; i.e., $f = \frac{h}{g}$ with $h, g \in F[\lambda_1, ..., \lambda_n]$ non-proportional monomials.

Remark 9.1.3. Whenever F is ν-archimedean, $\mathscr{L}$-monomials in $F(\lambda_1, ..., \lambda_n)$ are not bounded; i.e., for any $\mathscr{L}$-monomial f there does not exist $\alpha \in F$ for which $|f| \leq |\alpha|$.

Definition 9.1.4. A rational function $f \in F(\lambda_1, ..., \lambda_n)$ is a **hyperspace-fraction**, or HS-fraction, if $f \sim_{\langle F\rangle} \sum_{i=1}^t |f_i|$ where the f_i are non-proportional $\mathscr{L}$-monomials.

Remark 9.1.5. HS-fractions in $F(\lambda_1, ..., \lambda_n)$ are not bounded, since $|f_j| \leq \sum_{i=1}^t |f_i| = f$.

Definition 9.1.6. A **hyperplane ν-kernel**, or **HP-kernel**, for short, is is a principal ν-kernel of $F(\lambda_1, ..., \lambda_n)$ generated by an $\mathscr{L}$-monomial.

A **hyperspace-fraction ν-kernel**, or **HS-kernel**, for short, is is a principal ν-kernel of $F(\lambda_1, ..., \lambda_n)$ generated by a hyperspace fraction.

By definition, any HP-kernel is regular. Also, a fortiori, every HP-kernel is an HS-kernel.

Definition 9.1.7. $\Omega(\overline{F(\Lambda)})$ is the lattice of ν-kernels finitely generated by HP-kernels of $\overline{F(\Lambda)}$, i.e., every element $\langle f\rangle \in \Omega(\overline{F(\Lambda)})$ is obtained via finite intersections and products of HP-kernels.

Proposition 9.1.8. *Any principal HS-kernel is a product of distinct HP-kernels, and thus is in $\Omega(\overline{F(\Lambda)})$.*

Proof. If $\langle f\rangle$ is an HS-kernel, then $\langle f\rangle = \left\langle \sum_{i=1}^t |f_i| \right\rangle = \prod_{i=1}^t \langle f_i\rangle$ with $f_i \in F[\lambda_1, ..., \lambda_n]$ non-proportional $\mathscr{L}$-monomials. □

Corollary 9.1.9. *Any HS-Kernel is regular.*

Proof. Apply Proposition 8.0.11. □

Definition 9.1.10. A 1^ν-set in F^n is a **hyperplane 1^ν-set** (HP-1^ν-set for short) if it is defined by an $\mathscr{L}$-monomial. A 1^ν-set in F^n is a **hyperspace-fraction 1^ν-set** (HS-1^ν-set for short) if it is defined by an HS-fraction.

Corollary 9.1.11. *A 1^ν-set is an HS-1^ν-set if and only if it is an intersection of HP-1^ν-sets.*

Proof. As $1_{\mathrm{loc}}(\langle f\rangle\langle g\rangle) = 1_{\mathrm{loc}}(\langle f\rangle) \cap 1_{\mathrm{loc}}(\langle g\rangle)$, the assertion follows directly from Proposition 9.1.8. □

9.2. **Region ν-kernels.**

Our last kind of ν-kernel marks out a tropical region, and thus is called a "region ν-kernel." It also relates to information lost by passing from $F[\Lambda]$ to $\overline{F[\Lambda]}$.

Lemma 9.2.1. *Let $\langle f\rangle$ be an HP-kernel, with F divisible. If $w \in \langle f\rangle$ is an $\mathscr{L}$-monomial, then $w^s = f^k$ for some $s, k \in \mathbb{Z} \setminus \{0\}$.*

Proof. By assumption $\langle w\rangle \subseteq \langle f\rangle$, implying $1_{\mathrm{loc}}(w) \supseteq 1_{\mathrm{loc}}(f)$. Assume that $w^s \neq f^k$ for any $s, k \in \mathbb{Z} \setminus \{0\}$. We will show that there exists some $\mathbf{b} \in F^n$ such that $\mathbf{b} \in 1_{\mathrm{loc}}(f) \setminus 1_{\mathrm{loc}}(g)$ for some $g \in \langle f\rangle$, which clearly is impossible.

Let $(p_1, ..., p_n), (q_1, ..., q_n) \in \mathbb{Z}^n$ be the vectors of degrees of $\lambda_1, ..., \lambda_n$ in the Laurent monomials f and w. Since w and f are nonconstant, $(p_1, ..., p_n), (q_1, ..., q_n) \neq (0)$. Since F is divisible, we may take appropriate roots and assume that $gcd(p_1, ..., p_n) = gcd(q_1, ..., q_n) = 1$. Since $w \in \langle f\rangle$ we have $|w| \leq |f|^m$ for some $m \in \mathbb{N}$, so if λ_i occurs in w it must also occur in f. Finally, since $w^s \neq f^k$ for any $k \in \mathbb{Z} \setminus \{0\}$ we may also assume that $(p_1, ..., p_n) \neq (q_1, ..., q_n)$, for otherwise $w = \alpha f$ for some $\alpha \neq 1$ and thus

$$1_{\mathrm{loc}}(f) \subseteq 1_{\mathrm{loc}}(f) \cap 1_{\mathrm{loc}}(w) = 1_{\mathrm{loc}}(f) \cap 1_{\mathrm{loc}}(\alpha f) = \emptyset,$$

a contradiction. Suppose that ${\lambda_j}^l$ occurs in w for some $l \in \mathbb{Z} \setminus \{0\}$ such that λ_j is not identically 1 on $1_{\mathrm{loc}}(w)$ (and thus also on $1_{\mathrm{loc}}(f)$). Thus there exists some $k \in \mathbb{Z} \setminus \{0\}$ such that ${\lambda_j}^k$ occurs in f. Then λ_j does not occur in the Laurent monomial $g := w^{-k} f^{\ell} \in \langle f \rangle$. Without loss of generality, assume that $j = 1$. If

$$\mathbf{a} = (\alpha_1, ..., \alpha_n) \in 1_{\mathrm{loc}}(f),$$

then $g(\mathbf{a}) = w(\mathbf{a})^{-1} f(\mathbf{a}) = 1$. By assumption that $w^s \neq f^k$, there exists λ_t occurring in f and not in w. Take $\mathbf{b} = (1, \alpha_2, ..., \beta, ..., \alpha_n)$ with $\beta \in F$ occurring in the t-th component, such that $\beta^{p_t} = \frac{\alpha_t^{p_t}}{\alpha_1^{p_1}}$. Then as λ_j is not identically 1 over $1_{\mathrm{loc}}(f)$, we can choose $\mathbf{b} \in 1_{\mathrm{loc}}(w)$ such that $f(\mathbf{b}) = 1$ but $g(\mathbf{b}) \neq 1$. □

Proposition 9.2.2. *Let f be an $\mathscr{L}$-monomial. Then $w \in \langle f \rangle$ is an $\mathscr{L}$-monomial if and only if w is a generator of $\langle f \rangle$.*

Proof. The claim follows from Lemma 9.2.1 and the property that $\langle g^k \rangle = \langle g \rangle$. □

Definition 9.2.3. The **$\mathscr{L}$-binomial** o defined by an $\mathscr{L}$-monomial f is the rational function $1 + f$.

The **complementary $\mathscr{L}$-binomial** o^c of o is $1 + f^{-1}$. By definition $(\mathcal{O}^c)^c = \mathcal{O}$.

The **order ν-kernel** of the ν-semifield† $F(\lambda_1, ..., \lambda_n)$ defined by f is the principal ν-kernel $\mathcal{O} = \langle o \rangle$ for the $\mathscr{L}$-binomial $o = 1 + f$.

The **complementary order ν-kernel** $\mathcal{O}^c$ of $\mathcal{O}$ is $\langle o^c \rangle$.

Since f is an $\mathscr{L}$-monomial, so is f^{-1} and thus $\mathcal{O}^c$ is an order ν-kernel.

Lemma 9.2.4. *Let $\mathcal{O} = \langle 1 + f \rangle$ be an order ν-kernel of $F(\lambda_1, ..., \lambda_n)$. Then*

$$\mathcal{O} \cap \mathcal{O}^c = \langle 1 \rangle \text{ and } \mathcal{O}\mathcal{O}^c = \langle f \rangle.$$

Proof. $\mathcal{O} \cap \mathcal{O}^c = \langle |1+f| \wedge |1+f^{-1}| \rangle = \langle 1 \rangle$ and $\mathcal{O}\mathcal{O}^c = \langle |1+f| + |1+f^{-1}| \rangle = \langle 1 + f + f^{-1} \rangle = \langle 1 + |f| \rangle = \langle |f| \rangle = \langle f \rangle$ (noting that $(1 + f), (1 + f^{-1}) \geq_{\nu} 1$ implies $|1 + f| \cong_{\nu} 1 + f$ and $|1 + f^{-1}| \cong_{\nu} 1 + f^{-1}$). □

Definition 9.2.5. A rational function $f \in F(\lambda_1, ..., \lambda_n)$ is said to be a **region fraction** if $1_{\mathrm{loc}}(f)$ contains some nonempty open interval.

Lemma 9.2.6. *$f \sim \langle F \rangle \sum_{i=1}^{t} |o_i|$ is a region fraction iff, writing $o_i = 1 + f_i$ for $\mathscr{L}$-monomials f_i, we have $f_i \not\cong_{\nu} f_j^{\pm 1}$ for every $i \neq j$.*

Proof. Since $1 + f_i \geq_{\nu} 1$ for every i, we have that

$$\sum_{i=1}^{t} |o_i| = \sum_{i=1}^{t} |1 + f_i| = \sum_{i=1}^{t} (1 + f_i) = 1 + \sum_{i=1}^{t} f_i.$$

Thus a region fraction r can be defined as $r \sim_{\langle F \rangle} 1 + \sum_{i=1}^{t} f_i$, so the last condition of the definition can be stated as $f_j \neq f_i^{-1}$ for any $1 \leq i, j \leq t$. If there exist k and ℓ for which $f_{\ell} \cong_{\nu} f_k^{-1}$, then

$$\sum_{i=1}^{t} |o_i| \cong_{\nu} |f_k| + (1 + \sum_{i \neq k, \ell} f_i) \cong_{\nu} |f_k| + \Big|1 + \sum_{i \neq k, \ell} f_i\Big|,$$

and thus $1_{\mathrm{loc}}(r) = 1_{\mathrm{loc}}(\sum_{i=1}^{t} |o_i|) = 1_{\mathrm{loc}}(f_k) \cap 1_{\mathrm{loc}}(1 + \sum_{i \neq k, \ell} f_i) \subseteq 1_{\mathrm{loc}}(f_k)$.

□

Definition 9.2.7. A **region ν-kernel** is a principal ν-kernel generated by a region fraction.

Lemma 9.2.8. *A principal ν-kernel K is a region ν-kernel if and only if it has the form*

$$K = \prod_{i=1}^{v} \mathcal{O}_i$$

for order ν-kernels $\mathcal{O}_1, ..., \mathcal{O}_v \in \mathcal{P}(F(\lambda_1, ..., \lambda_n))$.

Proof. If $f = \sum |o_i|$ is a generating region fraction of K, then

$$K = \langle f \rangle = \langle \sum_{i=1}^{t} |o_i| \rangle \ = \ \prod_{j=1}^{v} \langle o_i \rangle.$$

Conversely, if $K = \prod_{i=1}^{v} \mathcal{O}_i$ then writing $\mathcal{O}_i = \langle o_i \rangle$, we see that $\sum_{i=1}^{t} |o_i|$ is a region fraction generating K, since $\left\langle \sum_{i=1}^{t} |o_i| \right\rangle = \prod_{i=1}^{v} \mathcal{O}_i = K$. □

Lemma 9.2.9. *Any order ν-kernel $f = 1 + \frac{h}{g}$ (for monomials g and h) is corner-internal.*

Proof. $f = \frac{g+h}{g}$. Any ghost root $\mathbf{a}$ of $\bar{f} = g^\nu + h$ is either a ghost root of g dominating h, or a ghost root of $g + h$; in either case it is a ν-kernel root of f. □

Lemma 9.2.10. *If $\langle f\rangle \neq \langle 1\rangle$ is a regular ν-kernel and $\langle o\rangle$ is an order ν-kernel, then $\langle f\rangle\langle o\rangle$ is regular.*

Proof. Write $f = \frac{h}{g}$ and $o = 1 + \frac{h'}{g'}$ with h' and g' monomials in $F[\lambda_1, ..., \lambda_n]$. Since regularity does not depend on the choice of generator of the ν-kernel and since

$$\langle f\rangle\langle o\rangle = \langle |f| + |o|\rangle = \langle |f| + o\rangle,$$

we check the condition on $|f| + o$. Then $|f| + o = |\frac{h}{g}| + 1 + \frac{h'}{g'}$. Since $|\frac{h}{g}| \geq_\nu 1$ we have

$$|f| + o = \left(\left|\frac{h}{g}\right| + 1\right) + \frac{h'}{g'} = \left|\frac{h}{g}\right| + \frac{h'}{g'} \cong_\nu \frac{(h^2 + g^2)g' + ghh'}{ghg'}.$$

Since $g \neq h$, we are done unless $g' >_\nu h'$, and Frobenius enables us to reduce to $\frac{h^2+g^2}{gh} = |f|$. But then we are done since $|f|$ is regular, by Lemma 7.4.3. □

Proposition 9.2.11. *$\langle f\rangle\langle o_1\rangle \cdots \langle o_k\rangle$ is regular, for any HP-kernel $\langle f\rangle \neq 1$ and order ν-kernels $\langle o_1\rangle, ..., \langle o_k\rangle$.*

Proof. Iterate Lemma 9.2.10, noting that every HP-kernel is regular. □

9.3. **Geometric interpretation of HS-kernels and region ν-kernels.**

Remark 9.3.1. In view of Theorem 8.2.1, for the case that F is a divisible, ν-archimedean supertropical ν-semifield†, we may take $F = \mathscr{R}$.

By Remark 3.4.5, any $\mathscr{L}$-monomial f may be considered as a linear functional over $\mathbb{Q}$ and thus the HP-kernel given by the equation $f = 1$ over $\mathscr{R}$ translates to $\mathfrak{f} = 0$ over $(\mathscr{R}, +)^{(n)}$ (in logarithmic notation), where $\mathfrak{f}$ is the linear functional obtained from f by applying (3.6).

Lemma 9.3.2. *If f is an $\mathscr{L}$-monomial in $F(\lambda_1, ..., \lambda_n)$, then f is completely determined by the set of vectors $\{w_0, ..., w_n\}$ for any $w_i = (\alpha_{i,1}, ..., \alpha_{i,n}, f(\alpha_{i,1}, ..., \alpha_{i,n})) \in F^{(n+1)}$ where $\{a_i = (\alpha_{i,1}, ..., \alpha_{i,n})\} \subset F^{(n)}$ such that $w_0, ..., w_n$ are in general position (are not contained in an $(n-1)$-dimensional affine subspace of $F^{(n+1)}$).*

Proof. Writing $f(\lambda_1, ..., \lambda_n) = \alpha \prod_{i=1}^n \lambda_i^{k_i}$ with $k_i \in \mathbb{Z}$, we have $\alpha = f(a_0) \prod_{i=1}^n a_{0,i}^{-k_i}$. After α is determined, since $w_0, ..., w_n$ are in general position the set

$$\{a_1, ..., a_n, b = (f(a_1), ..., f(a_n))\} \subset F^{(n)}$$

define a linearly independent set of n linear equations in the variables k_i, and thus determine them uniquely. □

Consider the HS-kernel of $F(\lambda_1, ..., \lambda_n)$ defined by the HS-fraction $f = \sum_{i=1}^t |f_i|$ where $f_1, ..., f_t$ are $\mathscr{L}$-monomials. Then $f(\mathbf{a}) \cong_\nu 1$ if and only if $f_i(\mathbf{a}) \cong_\nu 1$ for each $i = 1, ..., t$, which translates to a homogenous system of linear equations (over $\mathbb{Q}$) of the form $\mathfrak{f}_i = 0$ where $\mathfrak{f}_i$ is the logarithmic form of f_i. This way $1_{\mathrm{loc}}(f) \subset F^{(n)} \cong (F^+)^n$ is identified with an affine subspace of $F^{(n)}$ which is just the intersection of the t affine hyperplanes defined by $\mathfrak{f}_i = 0$, $1 \leq i \leq t$. Analogously, the $\mathscr{L}$-binomial $o = 1 + g$ has ν-value 1 if and only if $g \leq_\nu 1$, giving rise to the half space of $\mathbb{R}$ defined by the weak inequality $\mathfrak{g} \leq 0$. Thus, the region ν-kernel defined by $r = \sum_{i=1}^t |o_i| = 1 + \sum_{i=1}^t g_i$ where $o_i = 1 + g_i$ are $\mathscr{L}$-binomials, yields the nondegenerate polyhedron formed as the intersection of the affine half spaces each defined using the $\mathfrak{g}_i$ corresponding to o_i.

9.4. **The HO-decomposition.**

Since our next major goal, Theorem 9.4.8, is somewhat technical, we start this section with a motivating example. Any point $\mathbf{a} = (\alpha_1, ..., \alpha_n)$ in $F^{(n)}$ is just $1_{\mathrm{loc}}(f_\mathbf{a})$, where

$$f_\mathbf{a}(\lambda_1, ..., \lambda_n) = \left|\frac{\lambda_1}{\alpha_1}\right| + \cdots + \left|\frac{\lambda_n}{\alpha_n}\right| \in F(\lambda_1, ..., \lambda_n).$$

We would like $\langle f_\mathbf{a}\rangle$ to encode the reduction of dimension from $F^{(n)}$ to the point $\{\mathbf{a}\}$.

For each $1 \leq k \leq n$ define $f_{k,\mathbf{a}} = \left|\frac{\lambda_1}{\alpha_1}\right| + \cdots + \left|\frac{\lambda_k}{\alpha_k}\right|$ and $f_0 = 1$, and consider the chain of principal HS-kernels

$$\langle f_\mathbf{a}\rangle = \langle f_{n,\mathbf{a}}\rangle \supset \langle f_{n-1,\mathbf{a}}\rangle \supset \cdots \supset \langle f_{1,\mathbf{a}}\rangle \supset \{1\}. \tag{9.1}$$

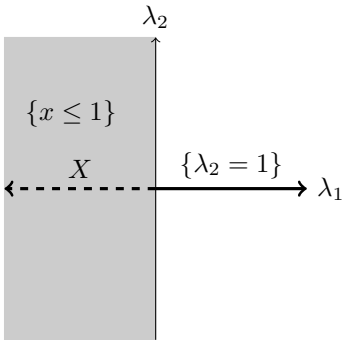

FIGURE 2. Order relations

The factors $\langle f_{k,\mathbf{a}}\rangle/\langle f_{k-1,\mathbf{a}}\rangle$ are the quotient semifields†

$$\prod_{j=1}^{k}\left\langle \frac{\lambda_j}{\alpha_j}\right\rangle \Big/ \prod_{j=1}^{k-1}\left\langle \frac{\lambda_j}{\alpha_j}\right\rangle \cong \left\langle \frac{\lambda_k}{\alpha_k}\right\rangle \Big/ \left(\left(\prod_{j=1}^{k-1}\left\langle \frac{\lambda_j}{\alpha_j}\right\rangle\right)\cap \left\langle \frac{\lambda_k}{\alpha_k}\right\rangle\right),$$

a nontrivial homomorphic image of an HP-kernel of the quotient ν-semifield† $F(\lambda_1,...,\lambda_n)\Big/\left(\left(\prod_{j=1}^{k-1}\langle\frac{\lambda_j}{\alpha_j}\rangle\right)\cap\langle\frac{\lambda_k}{\alpha_i}\rangle\right)$.

We claim that this chain of HS-kernels can be refined to a longer descending chain of principal ν-kernels descending from $\langle f_a\rangle$. Indeed, the ν-kernels $\langle|\lambda_1|+|\lambda_2|\rangle$ and $\langle|\lambda_1|\rangle = \langle\lambda_1\rangle$ both are semifields†, and $\langle|\lambda_1|\rangle$ is a ν-subkernel of $\langle|\lambda_1|+|\lambda_2|\rangle$. Consider the substitution map ϕ sending λ_1 to 1. Then $\operatorname{Im}(\phi) = \langle 1+|\lambda_2|\rangle = \langle|\lambda_2|\rangle_{F(\lambda_2)}$. The ν-kernel $\langle|\lambda_2|\rangle$ is not simple as a principal ν-kernel of the ν-semifield† $\langle|\lambda_2|\rangle_{F(\lambda_2)}$, for the chain $\langle|\lambda_2|\rangle \supset \langle|1+\lambda_2|\rangle \supset \langle 1\rangle$ is the image of the refinement

$$\langle|\lambda_1|+|\lambda_2|\rangle \supset \langle|\lambda_1+\lambda_2|+|\lambda_1|\rangle \supset \langle|\lambda_1|\rangle$$

(since $\phi(|\lambda_1+\lambda_2|+|\lambda_1|) = |\phi(\lambda_1)+\phi(\lambda_2)|+|\phi(\lambda_1)| = |1+\lambda_2|+|1| = |1+\lambda_2|$).

On the other hand, $\langle 1+\lambda_2\rangle$ is an order ν-kernel which induces the order relation $\lambda_2 \leq 1$ on the ν-semifield† $\langle|\lambda_2|\rangle$.

In view of these considerations, we would like to exclude order ν-kernels and ask:

- Can (9.1) be refined to a longer descending chain of HS-kernels descending from $\langle f_a\rangle$?
- Are the lengths of descending chains of HS-kernels beginning at $\langle f_a\rangle$ bounded?
- Can any chain of HS-kernels be refined to such a chain of maximal length?

We provide answers to these three questions, for which the chain (9.1) is of maximal unique length common to all chains of HS-kernels descending from $\langle f_a\rangle$. Our method is to provide an explicit decomposition of a principal ν-kernel $\langle f\rangle$ as an intersection of ν-kernels of two types: The first, called an **HO-kernel**, is a product of an HS-kernel and a region ν-kernel. The second is a product of a region ν-kernel and a bounded from below ν-kernel. Whereas the first type defines the 1^ν-set of $\langle f\rangle$, the second type corresponds to the empty set and thus has no effect on the geometry. This latter type is the source of ambiguity in relating a 1^ν-set to a ν-kernel, preventing the ν-kernel of the 1^ν-set from being principal.

When intersected with $\langle F\rangle$, the factors in the decomposition coming from ν-kernels of the second type are degenerate. Restriction to $\langle F\rangle$ thus removes our ambiguity, and each HO-kernel (intersected with $\langle F\rangle$) is in $1:1$ correspondence with its 1^ν-set (the segment in the 1^ν-set defined by $\langle f\rangle$). Then the 'HO-part' is unique and independent of the choice of the ν-kernel generating the 1^ν-set.

Geometrically, the decomposition to be described below is just the decomposition of a principal 1^ν-set defined by $\langle f\rangle$ to its "linear" components. Each component is obtained by bounding an affine subspace of $F^{(n)}$ defined by an appropriate HS-fraction (which in turn generates an HS-kernel) using a region fraction (generating a region ν-kernel). Although the HS-fraction and region fraction defining each segment may vary when we pass from one generator of a principal ν-kernel to the other, the HS-kernels and region ν-kernels are independent of the choice of generator, for they correspond to the components of the 1^ν-set of $\langle f\rangle$.

Construction 9.4.1. *Take a rational function $f \in F(\lambda_1,...,\lambda_n)$ for which $1_{\mathrm{loc}}(f) \neq \emptyset$. Replacing f by $|f|$, we may assume that $f \geq_\nu 1$. Write $f = \frac{h}{g} = \frac{\sum_{i=1}^k h_i}{\sum_{j=1}^m g_j}$ where h_i and g_j are monomials in $F[\lambda_1,...,\lambda_n]$. For each $\mathbf{a} \in 1_{\mathrm{loc}}(f)$, let*

$$H_a \subseteq H = \{h_i \ : \ 1 \leq i \leq k\}; \qquad G_a \subseteq G = \{g_j \ : \ 1 \leq j \leq m\}$$

be the sets of dominant monomials at $\mathbf{a}$; thus, $h_i(\mathbf{a}) = g_j(\mathbf{a})$ for any $h_i \in H_{\mathbf{a}}$ and $g_j \in G_{\mathbf{a}}$. Let $H_{\mathbf{a}}^c = H \setminus H_{\mathbf{a}}$ and $G_{\mathbf{a}}^c = G \setminus G_{\mathbf{a}}$. Then, for any $h' \in H_{\mathbf{a}}$ and $h'' \in H_{\mathbf{a}}^c$, $h'(\mathbf{a}) + h''(\mathbf{a}) = h'(\mathbf{a})$, or, equivalently, $1 + \frac{h''(\mathbf{a})}{h'(\mathbf{a})} = 1$. Similarly, for any $g' \in G_{\mathbf{a}}$ and $g'' \in G_{\mathbf{a}}^c$, $g'(\mathbf{a}) + g''(\mathbf{a}) = g'(\mathbf{a})$ or, equivalently, $1 + \frac{g''(\mathbf{a})}{g'(\mathbf{a})} = 1$.

Thus for any such $\mathbf{a}$ *we obtain the relations*

$$\frac{h'}{g'} = 1, \qquad \forall h' \in H_{\mathbf{a}},\ g' \in G_{\mathbf{a}}, \tag{9.2}$$

$$1 + \frac{h''}{h'} = 1;\ 1 + \frac{g''}{g'} = 1, \qquad \forall h' \in H_{\mathbf{a}},\ h'' \in H_{\mathbf{a}}^c,\ g' \in G_{\mathbf{a}},\ g'' \in G_{\mathbf{a}}^c. \tag{9.3}$$

As $\mathbf{a}$ *runs over* $1_{\mathrm{loc}}(f)$*, there are only finitely many possibilities for* H_a *and* G_a *and thus for the relations in* (9.2) *and* (9.3)*; we denote these as* $(\theta_1(i), \theta_2(i))$*,* $i = 1, ..., q$*.*

In other words, for any $1 \leq i \leq q$*, the pair* $(\theta_1(i), \theta_2(i))$ *corresponds to a* ν*-kernel* K_i *generated by the corresponding elements*

$$\frac{h'}{g'},\ \left(1 + \frac{h''}{h'}\right),\ \textit{and}\ \left(1 + \frac{g''}{g'}\right),$$

where $\{\frac{h'}{g'} = 1\} \in \theta_1$ *and* $\{1 + \frac{g''}{g'} = 1\},\ \{1 + \frac{h''}{h'} = 1\} \in \theta_2$*.*

Reversing the argument, every point satisfying one of these q *sets of relations is in* $1_{\mathrm{loc}}(f)$*. Hence,*

$$1_{\mathrm{loc}}\left(\langle f\rangle \cap \langle F\rangle\right) = 1_{\mathrm{loc}}(f) = \bigcup_{i=1}^{q} 1_{\mathrm{loc}}(K_i) = \bigcup_{i=1}^{q} 1_{\mathrm{loc}}\left(K_i \cap \langle F\rangle\right) \tag{9.4}$$

$$= 1_{\mathrm{loc}}\left(\bigcap_{i=1}^{q} (K_i \cap \langle F\rangle)\right),$$

Hence $\langle f\rangle \cap \langle F\rangle = \bigcap_{i=1}^{q} K_i \cap \langle F\rangle$*, since* $\langle f\rangle \cap \langle F\rangle,\ \bigcap_{i=1}^{q} K_i \cap \langle F\rangle \in \mathcal{P}(\langle F\rangle)$*.* $\bigcap_{i=1}^{q} K_i$ *provides a local description of* f *in a neighborhood of its* 1^ν*-set .*

Let us view this construction globally. We used the 1^ν*-set of* $\langle f\rangle$ *to construct* $\bigcap_{i=1}^{q} K_i$*. Adjoining various points* $\mathbf{a}$ *in* $F^{(n)}$ *might add some regions, complementary to the regions defined by* (9.3) *in* $\theta_2(i)$ *for* $i = 1, ..., q$*, over which* $\frac{h'}{g'} \neq 1, \forall h' \in H_{\mathbf{a}}, \forall g' \in G_{\mathbf{a}}$ *for each* $\mathbf{a}$*, i.e., regions over which the dominating monomials never agree. Continuing the construction above using* $\mathbf{a} \in F^{(n)} \setminus 1_{\mathrm{loc}}(f)$ *similarly produces a finite collection of, say* $t \in \mathbb{Z}_{\geq 0}$*,* ν*-kernels generated by elements from* (9.3) *and their complementary order fractions and by elements of the form* (9.2) *(where now* $\frac{h'}{g'} \neq 1$ *over the region considered). Any principal* ν*-kernel* $N_j = \langle q_j\rangle$*,* $1 \leq j \leq t$*, of this complementary set of* ν*-kernels has the property that* $1_{\mathrm{loc}}(N_j) = \emptyset$*, and thus by Corollary 8.1.11,* N_j *is bounded from below. As there are finitely many such* ν*-kernels there exists small enough* $\gamma >_\nu 1$ *in* $\mathcal{T}$ *for which* $|q_j| \wedge \gamma = \gamma$ *for* $j = 1, ..., t$*. Thus* $\bigcap_{j=1}^{t} N_j$ *is bounded from below and thus* $\bigcap_{j=1}^{t} N_j \supseteq \langle F\rangle$ *by Remark 8.1.9.*

Piecing this together with (9.4) *yields* f *over all of* $F^{(n)}$*, so we have*

$$\langle f\rangle = \bigcap_{i=1}^{q} K_i \cap \bigcap_{j=1}^{t} N_j. \tag{9.5}$$

So, $\langle f\rangle \cap \langle F\rangle = \bigcap_{i=1}^{q} K_i \cap \bigcap_{j=1}^{t} N_j \cap \langle F\rangle = \bigcap_{i=1}^{q} K_i \cap \langle F\rangle$*.*

In this way, we see that intersecting a principal ν-kernel $\langle f\rangle$ with $\langle F\rangle$ 'chops off' all of the bounded from below ν-kernels in (9.5) (the N_j's given above). This eliminates ambiguity in the ν-kernel corresponding to $1_{\mathrm{loc}}(f)$.
Finally we note that if $1_{\mathrm{loc}}(f) = \emptyset$, then $\langle f\rangle = \bigcap_{j=1}^{t} N_j$ for appropriate ν-kernels N_j and $\langle f\rangle \cap \langle F\rangle = \langle F\rangle$.

Remark 9.4.2.

(i) If K_1 and K_2 are such that $K_1 K_2 \cap F = \{1\}$ (i.e., $1_{\mathrm{loc}}(K_1) \cap 1_{\mathrm{loc}}(K_2) \neq \emptyset$), then the sets of $\mathscr{L}$-monomials θ_1 of K_1 and of K_2 are not the same (although one may contain the other), for otherwise together they would yield a single ν-kernel via Construction 9.4.1.

(ii) The ν-kernels K_i, being finitely generated, are in fact principal, so we can write $K_i = \langle k_i\rangle$ for rational functions $k_1, \ldots, k_q$. Let $\langle f\rangle \cap \langle F\rangle = \bigcap_{i=1}^{q} (K_i \cap \langle F\rangle) = \bigcap_{i=1}^{q} \langle |k_i| \wedge |\alpha|\rangle = \bigwedge_{i=1}^{q} \langle |k_i| \wedge |\alpha|\rangle$ with $\alpha \in F \setminus \{1\}$. By Theorem 8.5.3, for any generator f' of $\langle f\rangle \cap \langle F\rangle$ we have that $|f'| = \bigwedge_{i=1}^{q} |k_i'|$ with $k_i' \sim_{\langle F\rangle} |k_i| \wedge |\alpha|$ for every $i = 1, ..., q$. In particular, $1_{\mathrm{loc}}(k_i') = 1_{\mathrm{loc}}(|k_i| \wedge |\alpha|) = 1_{\mathrm{loc}}(k_i)$. Thus the ν-kernels K_i are independent of the choice of generator f, being defined by the components $1_{\mathrm{loc}}(k_i)$ of $1_{\mathrm{loc}}(f)$.

We now provide two instances of Construction 9.4.1. In what follows, we always make use of the above notation for the different types of ν-kernels involved in the construction. When denoting ν-kernels, R stands for "region", N for "null" (which are bounded from below), and L for "linear" (representing HS-kernels, which are unbounded).

Example 9.4.3. $F = (F, \mathcal{T}, \nu, \mathcal{G})$. Let $f = |\lambda_1| \wedge \alpha = \frac{\alpha|\lambda_1|}{\alpha + |\lambda_1|} \in F(\lambda_1, \lambda_2)$, where $\alpha >_\nu 1$ in $\mathcal{T}$. The order relation $\alpha \leq_\nu |\lambda_1|$ translates to the relation $\alpha + |\lambda_1| \cong_\nu |\lambda_1|$ or equivalently to $\alpha|\lambda_1|^{-1} + 1 \cong_\nu 1$. Over the region defined by this

relation we have $f = \frac{\alpha|\lambda_1|}{|\lambda_1|} = \alpha$. Similarly, its complementary order relation $\alpha \geq_\nu |\lambda_1|$ translates to $\alpha^{-1}|\lambda_1| + 1 = 1$ (via $|\lambda_1| + \alpha \cong_\nu \alpha$) over whose region $f = \frac{\alpha|\lambda_1|}{\alpha} = |\lambda_1|$. So

$$\langle f\rangle = K_1 \cap K_2 = (R_{1,1}L_{1,1}) \cap (R_{2,1}N_{2,1})$$

where $R_{1,1} = \langle \alpha^{-1}|\lambda_1|+1\rangle$, $L_{1,1} = \langle|\lambda_1|\rangle$, $R_{2,1} = \langle\alpha|\lambda_1|^{-1}+1\rangle$, and $N_{2,1} = \langle\alpha\rangle$. Geometrically $R_{1,1}$ is a strip containing the axis $\lambda_1 = 1$, and $R_{2,1}$ is the complementary region. The restriction of f to $R_{1,1}$ is $|\lambda_1|$ while f restricted to $R_{2,1}$ is α. Deleting $N_{2,1}$ we still have $1_{\mathrm{loc}}(f) = 1_{\mathrm{loc}}(R_{1,1}L_{1,1})$, although $R_{1,1}L_{1,1}$ properly contains $\langle f\rangle$.

$$\langle f\rangle = \langle f\rangle \cap \langle F\rangle = (R_{1,1}L_{1,1} \cap R_{2,1}N_{2,1}) \cap \langle F\rangle = (R_{1,1}L_{1,1}) \cap \langle F\rangle.$$

Example 9.4.4. Let $f = |\lambda_1+1|\wedge\alpha \in F(\lambda_1, \lambda_2)$ for some $\alpha >_\nu 1$ in $\mathcal{T}$. First note that $|\lambda_1+1| = \lambda_1+1$ since $\lambda_1+1 \geq 1$, allowing us to rewrite f as $(\lambda_1 + 1) \wedge \alpha$. Then $f = \frac{\alpha(\lambda_1+1)}{\alpha+(\lambda_1+1)} = \frac{\alpha\lambda_1+\alpha}{\alpha+\lambda_1}$. The order relation $\alpha \leq_\nu \lambda_1$ translates to the relation $\alpha+\lambda_1 \cong_\nu \lambda_1$ or equivalently to $\alpha\lambda_1^{-1}+1 \cong_\nu 1$, over whose region $f = \frac{\alpha\lambda_1+\alpha}{\alpha+\lambda_1} = \frac{\alpha\lambda_1+\alpha}{\lambda_1} = \alpha+\frac{\alpha}{\lambda_1} = \alpha$. Similarly, the complementary order relation $\alpha \geq \lambda_1$ translates to $\alpha^{-1}\lambda_1+1 \cong_\nu 1$ over whose region $f = \frac{\alpha\lambda_1+\alpha}{\alpha+\lambda_1} = \frac{\alpha\lambda_1+\alpha}{\alpha} = \lambda_1+1$. So

$$\langle f\rangle = K_1 \cap K_2 = (R_{1,1}\mathcal{O}_{1,2}) \cap (R_{2,1}N_{2,1}) = \mathcal{O}_{1,2} \cap (R_{2,1}N_{2,1})$$

where $R_{1,1} = \langle\alpha^{-1}\lambda_1+1\rangle$, $\mathcal{O}_{1,2} = \langle\lambda_1+1\rangle$, $R_{2,1} = \langle\alpha\lambda_1^{-1}+1\rangle$ and $N_{2,1} = \langle\alpha\rangle$. But $1_{\mathrm{loc}}(R_{2,1}N_{2,1}) \subseteq 1_{\mathrm{loc}}(N_{2,1}) = \emptyset$. So

$$1_{\mathrm{loc}}(f) = 1_{\mathrm{loc}}(\mathcal{O}_{1,2}) \cup 1_{\mathrm{loc}}(R_{2,1}N_{2,1}) = 1_{\mathrm{loc}}(\lambda_1 + 1) \cup \emptyset = 1_{\mathrm{loc}}(\lambda_1 + 1).$$

Now suppose L_i are HP-kernels and $\mathcal{O}_j$ are order ν-kernels, and let $L = \prod L_i$ and $R = \prod \mathcal{O}_j$. As can be seen easily from examples 9.4.3 and 9.4.4, by substituting any $\mathscr{L}$-monomial for λ_1 and any order fraction for $\lambda_1 + 1$,

$$\begin{aligned}(LR) \cap \langle F\rangle &= \prod(L_i \cap \langle F\rangle) \prod(\mathcal{O}_j \cap \langle F\rangle)\\ &= \Big(\prod((L_iR_i) \cap (N_iR_i^c)) \prod\big((\mathcal{O}_jR'_j) \cap (M_jR'^{\,c}_j)\big)\Big) \cap \langle F\rangle\\ &= \Big(\prod L_i \prod \mathcal{O}'_j\Big) \cap N \cap \langle F\rangle\end{aligned}$$

where the R_i and R'_j are region ν-kernels and $N = \left(\prod_i N_i\right)\left(\prod_j M_j\right)$. The ν-kernels N_i, M_j, and thus N, are bounded from below; $\langle F\rangle \subseteq N$, and the right side is $(LR') \cap \langle F\rangle$ where R' is a region ν-kernel.

Note that the $\mathcal{O}'_j$s involve the R_is, the R'_js and the $\mathcal{O}_j$s. Also note that intersecting with $\langle F\rangle$ keeps the HS-kernel unchanged in the new decomposition.

As the N_js in (9.5), being bounded from below, do not affect $1_{\mathrm{loc}}(f)$, we put them aside for the time being and proceed to study the structure of the ν-kernels K_i and their corresponding 1^ν-sets.

Take one of these ν-kernels K_i. Recall that K_i is generated by a set comprised of $\mathscr{L}$-monomials and order elements. Let $L_{i,j}$, $1 \leq j \leq u$, and $\mathcal{O}_{i,k}$, $1 \leq k \leq v$, be the HP-kernels and the order ν-kernels generated respectively by these $\mathscr{L}$-monomials and order elements. Then we can write

$$K_i = L_iR_i = \left(\prod_{j=1}^{u} L_{i,j}\right)\left(\prod_{k=1}^{v} \mathcal{O}_{i,k}\right), \tag{9.6}$$

where $L_i = \prod_{j=1}^u L_{i,j}$ is an HS-kernel and $R_i = \prod_{k=1}^v \mathcal{O}_{i,k}$ is a region ν-kernel. By assumption, $1_{\mathrm{loc}}(K_i) \neq \emptyset$ since at least one point of the 1^ν-set was used in its construction. Moreover, one cannot write $L = M_1 \cap M_2$ for distinct HS-kernels M_1 and M_2, for otherwise the construction would have produced two distinct ν-kernels, one with M_1 as its HS-kernel and the other with M_2 as its HS-kernel, rather than K_i.

Let us formalize this situation.

Definition 9.4.5. A rational function $f \in F(\lambda_1, ..., \lambda_n)$ is an **HO-fraction** if it is the sum of an HS-fraction f' and a region fraction o_f.

Definition 9.4.6. A principal ν-kernel $K \in \mathcal{P}(F(\lambda_1, ..., \lambda_n))$ is said to be an **HO-kernel** if it is generated by an HO-fraction.

Note that any HS-kernel or any region ν-kernel is an HO-kernel.

Lemma 9.4.7. *A principal ν-kernel K is an HO-kernel if and only if $K = LR$ where L is an HS-kernel and R is a region ν-kernel.*

Proof. ($\Rightarrow$) Write $K = \langle f\rangle$, where $f = f' + o_f$ is an HO-fraction. Thus $K = \langle f' + o_f\rangle = \langle f'\rangle\langle o_f\rangle$ where $\langle f'\rangle$ is an HS-kernel and $\langle o_f\rangle$ is a region ν-kernel.

($\Leftarrow$) Write the HO-fraction $f = f'+r$ where f' is an HS-fraction generating L and r is a region fraction generating R; then $\langle f\rangle = \langle f' + r\rangle = \langle f'\rangle\langle r\rangle = LR = K$. □

Theorem 9.4.8. *Every principal ν-kernel $\langle f\rangle$ of $F(\lambda_1, ..., \lambda_n)$ can be written as the intersection of finitely many principal ν-kernels*

$$\{K_i : i = 1, ..., q\} \text{ and } \{N_j : j = 1, ..., m\},$$

whereas each K_i is the product of an HS-kernel and a region ν-kernel

$$K_i = L_i R_i = \prod_{j=1}^{t_i} L_{i,j} \prod_{k=1}^{k_i} \mathcal{O}_{i,k} \tag{9.7}$$

while each N_j is a product of bounded from below ν-kernels and (complementary) region ν-kernels. For $\langle f\rangle \in \mathcal{P}(\langle F\rangle)$, the N_j can be replaced by $\langle F\rangle$ without affecting the resulting ν-kernel.

Proof. Let $K = LR$ be an HO-kernel with R a region ν-kernel and L an HS-kernel. By Proposition 9.1.8 and Lemma 9.2.8, we have that $L = \prod_{i=1}^{u} L_i$ for some HP-kernels $L_1, ..., L_u$ and $R = \prod_{j=1}^{v} \mathcal{O}_j$ for some order ν-kernels $\mathcal{O}_1, ..., \mathcal{O}_v$. Thus

$$K = LR = \left(\prod_{i=1}^{u} L_i\right)\left(\prod_{j=1}^{v} \mathcal{O}_j\right), \tag{9.8}$$

where $u \in \mathbb{Z}_{\geq 0}$, $v \in \mathbb{N}$, $L_1, ..., L_u$ are HP-kernels and $\mathcal{O}_1, ..., \mathcal{O}_v$ are order ν-kernels.

Let K_1 and K_2 be region ν-kernels (respectively HS-kernels) such that $(K_1K_2) \cap F = \{1\}$. Then K_1K_2 is a region ν-kernel (respectively HS-kernel). Consequently, if K_1 and K_2 are HO-kernels such that $(K_1K_2) \cap F = \{1\}$, then K_1K_2 is an HO-kernel. Indeed, the assertions follow from the decomposition $K_i = L_i\mathcal{O}_i = (\prod_{j=1}^{u_i} L_{i,j})(\prod_{k=1}^{v_i} \mathcal{O}_{i,k})$ for $i = 1, 2$, so that

$$K_1K_2 = (L_1L_2)(\mathcal{O}_1\mathcal{O}_2) = \left(\prod_{i=1}^{u_1} L_{1,i} \prod_{i=1}^{u_2} L_{2,i}\right)\left(\prod_{j=1}^{v_1} \mathcal{O}_{1,j} \prod_{j=1}^{v_2} \mathcal{O}_{2,j}\right) = L\mathcal{O},$$

with the appropriate u_i, v_i taken for $i = 1, 2$. □

Proposition 9.4.9.

- *If $\langle f\rangle$ is an HS-kernel, then the decomposition degenerates to $\langle f\rangle = K_1$ with $K_1 = L_1 = \langle f\rangle$.*
- *If $\langle f\rangle$ is a region ν-kernel, then $\langle f\rangle = K_1$ with $K_1 = R_1 = \langle f\rangle$.*
- *$\langle f\rangle$ is an irregular ν-kernel if and only if there exists some $i_0 \in \{1, ..., q\}$ such that $K_{i_0} = R_{i_0} = \prod_{k=1}^{k} \mathcal{O}_{k,i_0}$.*
- *$\langle f\rangle$ is a regular ν-kernel if and only if K_i is comprised of at least one HP-kernel, for every $i = 1, ..., q$.*

Proof. (i) By Proposition 9.1.8, $\langle f\rangle = \prod_{j=1}^{t} L_j$ where L_j is an HP-kernel for each j. Write $f = \frac{\sum h_i}{\sum g_j}$ for monomials $h_i, g_j \in F[\Lambda]$. If $\langle f\rangle$ is irregular, then at some neighborhood of a point $\mathbf{a} \in F^{(n)}$ we have some i_0 and j_0 for which $h_{i_0} \cong_\nu g_{j_0}$ where $(g_{j_0}(\mathbf{a}) \cong_\nu) h_{i_0}(\mathbf{a}) >_\nu h_i(\mathbf{a}), g_j(\mathbf{a})$ for every $i \neq i_0$ and $j \neq j_0$. The ν-kernel corresponding to (the closure) of this region has its relations (9.2) degenerating to $1 = 1$ as $\frac{h_{i_0}}{g_{j_0}} \cong_\nu 1$ over the region, thus is given only by its order relations of (9.3).

The last three assertions are direct consequences of (9.7); namely, if $\langle f\rangle$ is either an HS-kernel or a region ν-kernel, $\langle f\rangle$ already takes on the form of its decomposition. The fourth is equivalent to the third.

□

9.5. **Conclusion of the proof of Theorem 8.0.13.**

Having the palate of ν-kernels at our disposal, we can now better understand Theorem 8.0.13.

Definition 9.5.1. The **HO-decomposition** of a principal ν-kernel $\langle f\rangle$ is its decomposition given in Theorem 9.4.8. In the special case where $\langle f\rangle \in \mathcal{P}(\langle F\rangle)$, all bounded from below terms of the intersection are equal to $\langle F\rangle$.

Definition 9.5.2. For a subset $S \subseteq F(\lambda_1, ..., \lambda_n)$, denote by $\mathrm{HO}(S)$ the family of HO-fractions in S, by $\mathrm{HS}(S)$ the family of HS-fractions in S, and by $\mathrm{HP}(S)$ the family of $\mathscr{L}$-monomials in S.

Remark 9.5.3. $\mathrm{HP}(S) \subset \mathrm{HS}(S) \subset \mathrm{HO}(S)$ for any $S \subseteq F(\lambda_1, ..., \lambda_n)$, since every $\mathscr{L}$-monomial is an HS-fraction and every HS-fraction is an HO-fraction.

Example 9.5.4. Consider the ν-kernel $\langle f\rangle$ where $f = \frac{\lambda_1}{\lambda_2+1} \in F(\lambda_1, ..., \lambda_n)$. The points on the 1^ν-set of f define three distinct HS-kernels: $\langle \frac{\lambda_1}{\lambda_2}\rangle$ (corresponding to $\lambda_1 = \lambda_2$) over the region $\{\lambda_2 \geq 1\}$ defined by the region ν-kernel $\langle 1 + \lambda_2^{-1}\rangle$, $\langle \lambda_1\rangle = \langle \frac{\lambda_1}{1}\rangle$ (corresponding to $\lambda_1 = 1$) over the region $\{\lambda_2 \leq 1\}$ defined by the region ν-kernel $\langle 1 + \lambda_2\rangle$, and $\langle |\lambda_1| + |\lambda_2|\rangle$ (corresponding to the point defined by $\lambda_1 = \lambda_2 = 1$). Thus by Construction 9.4.1,

$$\langle f\rangle = \left(\left\langle \frac{\lambda_1}{\lambda_2} \right\rangle \langle 1 + \lambda_2^{-1}\rangle\right) \cap \langle \lambda_1\rangle\langle 1 + \lambda_2\rangle \cap \langle |\lambda_1| + |\lambda_2|\rangle\langle 1\rangle$$

and

$$1_{\mathrm{loc}}(f) = \left(1_{\mathrm{loc}}\left(\frac{\lambda_1}{\lambda_2}\right) \cap 1_{\mathrm{loc}}(1 + \lambda_2^{-1})\right) \cup (1_{\mathrm{loc}}(\lambda_1) \cap 1_{\mathrm{loc}}(1 + \lambda_2)) \cup \left(1_{\mathrm{loc}}(|\lambda_1| + |\lambda_2|) \cap F^{(2)}\right).$$

The third component of the decomposition (i.e., the HS-kernel $\langle |\lambda_1| + |\lambda_2|\rangle$) could be omitted without effecting $1_{\mathrm{loc}}(f)$.

The decomposition is shown (in logarithmic scale) in Figure 3, where the first two components are the rays emanating from the origin and the third component is the origin itself.

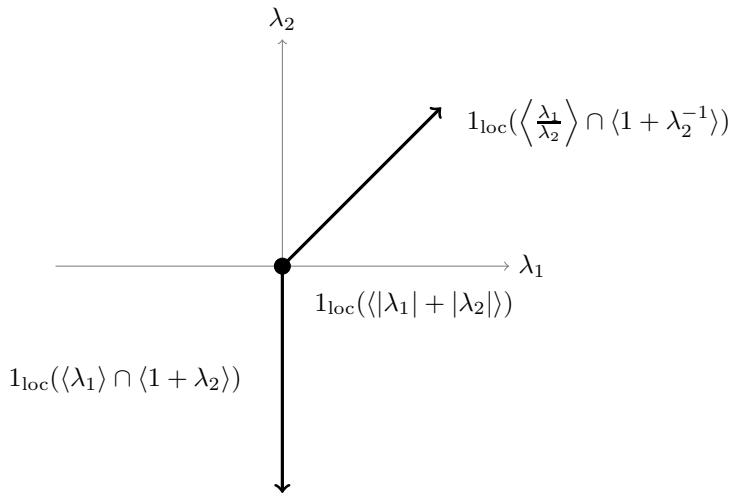

FIGURE 3. $1_{\mathrm{loc}}(\frac{\lambda_1}{\lambda_2+1}) = 1_{\mathrm{loc}}(\left\langle \frac{\lambda_1}{\lambda_2} \right\rangle \cap \langle 1+\lambda_2^{-1}\rangle) \cup 1_{\mathrm{loc}}(\langle \lambda_1 \rangle \cap \langle 1+\lambda_2 \rangle) \cup 1_{\mathrm{loc}}(|\lambda_1|+|\lambda_2|)$

9.5.1. *The lattice generated by regular corner-internal principal ν-kernels.*

Recall from Proposition 9.2.11 that the principal ν-kernel $\langle f\rangle\langle o_1\rangle\cdots\langle o_k\rangle$ is regular, for any HP-kernel $\langle f\rangle \neq 1$ and order ν-kernels $\langle o_1\rangle, ..., \langle o_k\rangle$.

Lemma 9.5.5. *Let $K \in \mathcal{P}(\langle F\rangle)$ and let*

$$K = (LR) \cap \langle F\rangle = \left(\prod_{j=1}^{u} L_j\right)\left(\prod_{k=1}^{v} \mathcal{O}_k\right) \cap \langle F\rangle$$

be the decomposition of K, as given in (9.6)*, where $L_1, ..., L_u$ are HP-kernels and $\mathcal{O}_1, ..., \mathcal{O}_v$ are order ν-kernels. If $u \neq 0$, i.e., $L \neq \langle 1\rangle$, then K is regular.*

Proof. Indeed, $K = (\prod_{j=2}^{u} L_j)(L_1 \prod_{k=1}^{v} \mathcal{O}_k) \cap \langle F\rangle$. By Proposition 9.2.11, $(L_1 \prod_{j=1}^{v} \mathcal{O}_j)$ is regular since L_1 is an HP-kernel. Thus K is a regular ν-kernel, since a product of regular ν-kernels is regular and since intersection with $\langle F\rangle$ does not affect regularity. □

We are finally ready to prove Theorem 8.0.13.

Proof. Let $\langle f\rangle$ be a principal ν-kernel and let

$$\langle f\rangle = \left(\bigcap_{i=1}^{q} K_i\right) \cap \langle F\rangle, \quad K_i = \prod_{j=1}^{t} L_{i,j} \prod_{k=1}^{k} \mathcal{O}_{i,k}$$

be its HO-decomposition. By Lemma 9.2.9, each HP-kernel $L_{i,j}$ and each order ν-kernel $\mathcal{O}_{i,k}$ are corner internal. Thus $\langle f\rangle$ as a finite product of principal corner internal ν-kernels is in the lattice generated by principal corner-internal ν-kernels.

For the second assertion, if $\langle f\rangle$ is regular, then by Theorem 9.4.8, for every $1 \leq i \leq q$, we have that $L_{i,1} \neq 1$. Thus by Lemma 9.5.5, each K_i is a product of principal regular corner-internal ν-kernels. Thus $\langle f\rangle$ is in the lattice generated by principal regular corner-internal ν-kernels. We conclude with Proposition 8.0.11. □

10. Polars: an intrinsic description of $\mathcal{K}$-kernels

To characterize $\mathcal{K}$-kernels intrinsically, we need a kind of orthogonality relationship, and introduce a new kind of ν-kernel, called **polar**, borrowed from the theory of lattice ordered groups ([36, section (2.2)]). We fix $X \subseteq F^{(n)}$. $\mathcal{S}$ is always assumed to be an idempotent 1-ν-semifield†, often $F(X)$ or even $\mathrm{Fun}(F^{(n)}, F)$. In this section, we lay out the general basics of the theory, although its full strength is only obtained in the following sections when the underlying ν-semifield† F is taken to be ν-divisible, ν-archimedean, and complete, e.g., $F = \mathscr{R}$.

10.1. Basic properties of polars.

Definition 10.1.1. We write $f \perp g$ for $f, g \in F(X)$ if $|f(\mathbf{a})| \wedge |g(\mathbf{a})| \cong_\nu 1$ for all $\mathbf{a} \in X$, i.e., $1_{\mathrm{loc}}(f) \cup 1_{\mathrm{loc}}(g) = X$.

For subsets K, L of $F(X)$ we write $K \perp L$ if $f \perp g$ for all $f \in K$ and $g \in L$. (This is not necessarily implied by $(\wedge|K|) \perp (\wedge|L|)$, as evidenced by Note 5.2.8.) For $V \subseteq \mathcal{S}$, we define

$$V^\perp = \{g \in \mathcal{S} : g \perp V.\} \tag{10.1}$$

Such a set $V^\perp$ is called a **polar** in the literature.

For $f \in \mathcal{S}$ we write $f^\perp$ for $\{f\}^\perp$. Thus, $f^\perp = |f|^\perp$. The set of all polars in $\mathcal{S}$ is denoted as $\mathcal{P}lr(\mathcal{S})$.

Remark 10.1.2. $f \perp g$ iff $|f| \perp g$, iff $\langle f\rangle \perp \langle g\rangle$.

Remark 10.1.3 ([36]). If $K \subset \mathcal{S}$, then $K^\perp = \langle K\rangle^\perp$. Consequently,

$$\langle K\rangle^{\perp\perp} = K^{\perp\perp}.$$

Thus, $L \perp K$ iff $K \subseteq L^\perp$, for any kernels K, L.

Although this kind of property cannot arise in classical algebraic geometry since a variety cannot be the proper union of two algebraic sets, it is quite common in the tropical setting. For example $(1+f) \perp (1+f^{-1})$. The usual properties of orthogonality go through here.

Lemma 10.1.4. *The following statements are immediate consequences of Definition 10.1.1. For any $K, L \subset \mathcal{S}$,*

(i) $K \subseteq L \Rightarrow K^\perp \supseteq L^\perp$.
(ii) $K \subseteq K^{\perp\perp}$.
(iii) $K^\perp = K^{\perp\perp\perp}$.
(iv) *K is a polar iff $K^{\perp\perp} = K$.*

Proof. The first two assertions are obvious, and the third follows from using $K^\perp$ in (ii), applying (i). Finally, if K is a polar, then $K = V^\perp$ for some $V \subseteq \mathcal{S}$, and $K^{\perp\perp} = (V^\perp)^{\perp\perp} = V^\perp = K$. □

The following facts, taken from [36], can be checked pointwise.

Theorem 10.1.5. *For any subset V of $\mathcal{S}$, $V^\perp$ is a $\mathcal{K}$-kernel of $\mathcal{S}$.*

Proof. As noted in [36, Theorem 2.2.4(e)], $K := V^\perp$ is a convex (abelian) subgroup, and thus a ν-kernel. It remains to show that $K = \mathcal{K}ern(1_{\mathrm{loc}}(K))$. Clearly $\subseteq$ holds, so we need to show that any $g \in \mathcal{K}ern(1_{\mathrm{loc}}(K))$ belongs to K. We are given $1_{\mathrm{loc}}(g) \supseteq 1_{\mathrm{loc}}(K)$ and $K \perp v$ for each v in V, so $1_{\mathrm{loc}}(g) \supseteq 1_{\mathrm{loc}}(K)$ implies that $1_{\mathrm{loc}}(g) \cup 1_{\mathrm{loc}}(v) \supseteq 1_{\mathrm{loc}}(K) \cup 1_{\mathrm{loc}}(v) = X$, i.e., $g \perp V$, as desired. □

Proposition 10.1.6. *$(\mathcal{P}lr(\mathcal{S}), \cdot, \cap, \perp, \{1\}, \mathcal{S})$ is a complete Boolean algebra.*

Proof. [36, Theorem 2.2.5]; If $\{K_i; i \in I\}$ are subsets of $\mathcal{S}$, then

$$\left(\bigcup_i K_i\right)^\perp = \bigcap_{i\in I} K_i^\perp.$$

□

Closure under complements is a consequence of (ii).

Proposition 10.1.7. *For any subset $V \subseteq \mathcal{S}$, $V^{\perp\perp}$ is the minimal polar containing V.*

Proof. By definition, $V^{\perp\perp}$ is a polar containing S. Let $P \supseteq S$ be a polar. Then $S^{\perp\perp} = (S^\perp)^\perp \subseteq (P^\perp)^\perp = P$. □

Definition 10.1.8. Let $S \subseteq \mathcal{S}$. We say that a polar P is **generated** by S if $P = S^{\perp\perp}$. If $S = \{f\}$ then we also write $f^{\perp\perp}$ for the polar generated by $\{f\}$.

Definition 10.1.9. A polar P of $\mathcal{S}$ is **principal** if there exists some $f \in \mathcal{S}$ such that $P = f^{\perp\perp} (= \langle f\rangle^{\perp\perp})$.

Lemma 10.1.10. *The collection of principal polars is a sublattice of $(\mathcal{P}lr(\mathcal{S}), \cdot, \cap)$.*

Proof. For any $f, g \in \mathcal{S}$,

$$(\langle f\rangle \cap \langle g\rangle)^{\perp\perp} = (\langle |f| \wedge |g|\rangle)^{\perp\perp} = (|f| \wedge |g|)^{\perp\perp},$$

and

$$(\langle f\rangle\langle g\rangle)^{\perp\perp} = (\langle |f| + |g|\rangle)^{\perp\perp} = (|f| + |g|)^{\perp\perp},$$

by Corollary 4.4. □

Remark 10.1.11. Any ν-kernel K of an idempotent ν-semifield† $\mathcal{S}$ cannot be orthogonal to a nontrivial ν-subkernel of $K^{\perp\perp}$.

Proof. Any ν-kernel L of $K^{\perp\perp}$ is a ν-kernel of $\mathcal{S}$. If $K \perp L$, then $L \subseteq K^\perp \cap K^{\perp\perp} = \{1\}$, yielding $L = \{1\}$. □

10.2. **Polars over complete Archimedean semifields†.**

In Theorem 5.3.6 we showed that $\mathcal{K}$-kernels of $\overline{F(\Lambda)}$ correspond to 1^ν-sets in $F^{(n)}$. Our aim here is to characterize these special kind of ν-kernels as polars, as a converse to Theorem 10.1.5. In analogy to algebraic geometry, polars play the role of 'radical ideals,' corresponding to root sets by an analog to the celebrated Nullstellensatz theorem, but for this we need to make extra assumptions on the underlying ν-semifield†.

10.3. **The polar-1^ν-set correspondence.**

Recall that $\mathscr{R}$ is presumed to be a ν-bipotent, divisible, archimedean and complete ν-semifield†. In this subsection we concentrate on $\overline{\mathscr{R}(\Lambda)}$. Doing so, we consider the natural extensions to $\overline{\mathscr{R}(\Lambda)}$ of the operators 1_{loc} and $\mathcal{K}ern$ defined in §5.3 with respect to $\mathscr{R}(\Lambda)$. We write $\mathcal{K}ern_{\mathscr{R}(\Lambda)}$ to denote the restriction of $\mathcal{K}ern$ to $\mathscr{R}(\Lambda)$.

Recall "completely closed" from Definition 2.2.8.

Lemma 10.3.1. *All ν-kernels of $\overline{\mathscr{R}(\Lambda)}$ that are completely closed are completions of ν-kernels of $\mathscr{R}(\Lambda)$.*

Proof. By Theorem 4.1.12(1), the ν-kernels of $\mathscr{R}(\Lambda)$ are precisely $\{K \cap \mathscr{R}(\Lambda) : K \text{ is a } \nu\text{-kernel of } \overline{\mathscr{R}(\Lambda)}\}$. Since $\mathscr{R}(\Lambda)$ is dense in $\overline{\mathscr{R}(\Lambda)}$ for every ν-kernel K of $\overline{\mathscr{R}(\Lambda)}$, the ν-kernel $L = K \cap \mathscr{R}(\Lambda)$ of $\mathscr{R}(\Lambda)$ is dense in K. Since $L \subseteq K$, we conclude that $\bar{L} = \bar{K}$, i.e., $\bar{L} = K$ if and only if K is completely closed. □

Proposition 10.3.2. *Each completely closed ν-kernel K of $\overline{\mathscr{R}(\Lambda)}$ defines a unique ν-kernel of $\mathscr{R}(\Lambda)$ given by $L = K \cap \mathscr{R}(\Lambda)$ for which $\overline{L} = K$.*

Example 10.3.3. Consider the ν-kernel $K = \langle |\lambda_1| \wedge |\alpha| \rangle \in \mathcal{P}(\mathscr{R}(\lambda_1))$ and its subset $X = \{|\lambda_1| \wedge |\alpha|^n : n \in \mathbb{N}\}$. Then $|\lambda_1| = \bigvee_{f \in X} f \in \overline{K}$. Hence $\langle \lambda_1 \rangle \subset \overline{K}$ yielding $\overline{\langle |\lambda_1| \wedge |\alpha| \rangle} = \overline{\langle |\lambda_1| \rangle}$.

Remark 10.3.4. $1_{\mathrm{loc}}(\bar{K}) = 1_{\mathrm{loc}}(K)$ for every ν-kernel K of $\mathscr{R}(\Lambda)$.

Proof. First note that $1_{\mathrm{loc}}(\bar{K}) \subseteq 1_{\mathrm{loc}}(K)$ since $K \subset \bar{K}$ in $\overline{\mathscr{R}(\Lambda)}$. Now, let $X = 1_{\mathrm{loc}}(K)$. For any nonempty subset A of K. If $\bigvee_{f \in A} f \in \overline{\mathscr{R}(\Lambda)}$ then for any $a \in X$,

$$\left(\bigvee_{f \in A} f \right)(\mathbf{a}) = \bigvee_{f \in A} f(\mathbf{a}) = \bigvee_{f \in A} 1 \cong_\nu 1,$$

yielding $1_{\mathrm{loc}}(\bigvee_{f \in A}) \supseteq X$ (cf. Lemma 8.1.13). Similarly, if $\bigwedge_{f \in A} f \in \overline{\mathscr{R}(\Lambda)}$ then for any $\mathbf{a} \in X$ we have $(\bigwedge_{f \in A} f)(\mathbf{a}) = \bigwedge_{f \in A} f(\mathbf{a}) = \bigwedge_{f \in A} 1 \cong_\nu 1$, yielding that $1_{\mathrm{loc}}(\bigwedge_{f \in A} f) \supseteq X$. We conclude that $1_{\mathrm{loc}}(\bar{K}) = 1_{\mathrm{loc}}(K)$. □

Proposition 10.3.5. *Every $\mathcal{K}$-kernel of $\overline{\mathscr{R}(\Lambda)}$ is completely closed.*

Proof. By Lemma 5.3.5, if $Z = 1_{\mathrm{loc}}(K)$ then $1_{\mathrm{loc}}(\mathcal{K}ern(Z)) = Z$. Let $Z = 1_{\mathrm{loc}}(K) \subseteq \mathscr{R}^{(n)}$ where $K = \mathcal{K}ern(Z)$, a ν-kernel of $\overline{\mathscr{R}(\Lambda)}$. $1_{\mathrm{loc}}(f) \supseteq Z$ for any $f \in K$. Now, if $\bigvee_{f \in A} f \in \overline{\mathscr{R}(\Lambda)}$ and $\bigwedge_{f \in A} f \in \overline{\mathscr{R}(\Lambda)}$, then by Lemma 8.1.13, for any $a \in Z$

$$\left(\bigvee_{f \in A} f \right)(a) = \bigvee_{f \in A} f(a) = \bigvee_{f \in A} 1 \cong_\nu 1 \quad \text{and} \quad (\bigwedge_{f \in A} f)(a) = \bigwedge_{f \in A} f(a) = \bigwedge_{f \in A} 1 \cong_\nu 1.$$

Thus

$$1_{\mathrm{loc}} \left(\bigvee_{f \in A} f \right), 1_{\mathrm{loc}} \left(\bigwedge_{f \in A} f \right) \supseteq Z.$$

So $\bigvee_{f \in A} f, \bigwedge_{f \in A} f \in K$, and K is completely closed in $\overline{\mathscr{R}(\Lambda)}$. □

Remark 10.3.6. Proposition 10.3.5 is not true when taking $\mathcal{K}$-kernels of $\mathscr{R}(\Lambda)$ instead of $\overline{\mathscr{R}(\Lambda)}$. Let $\alpha \in \mathscr{R}$ such that $\alpha > 1$. Consider the subset

$$X = \{|\lambda_1^n| \wedge \alpha : n \in \mathbb{N}\}.$$

Then $X \subset \langle \lambda_1 \rangle$, $\bigvee_{f \in X} f = \alpha$ (the constant function) and $1_{\mathrm{loc}}(f) = \{1\} \subset \mathscr{R}$ for every $f \in X$. Thus $\alpha = (\alpha)(1) = (\bigvee_{f \in X} f)(1) \neq \bigvee_{f \in X} f(1) = 1$ and α is not in the preimage of $1_{\mathrm{loc}}(\lambda_1)$. So we deduce that $\mathcal{K}$-kernels of $\mathscr{R}(\Lambda)$ need not be completely closed in $\mathscr{R}(\Lambda)$, and thus not polars since every polar is completely closed by Proposition 10.3.5. Also note that $\alpha \in \lambda_1^{\perp\perp}$, yielding $\lambda_1^\perp = \lambda_1^{\perp\perp\perp} = (\lambda_1^{\perp\perp})^\perp = \{1\}$.

Theorem 10.3.7. *The following properties of a ν-kernel K of $\overline{\mathscr{R}(\Lambda)}$ are equivalent:*

(i) *K is a $\mathcal{K}$-kernel.*
(ii) *K is a polar.*
(iii) *K is completely closed.*

Proof. *(i)* ⇒ *(iii)* By Proposition 10.3.5. $\mathcal{K}$-kernel of $\overline{\mathscr{R}(\Lambda)}$ is completely closed.
(iii) ⇒ *(ii)* See [36, Theorem 2.3.7]. One checks the condition of Lemma 10.1.4.
(ii) ⇒ *(i)* By Theorem 10.1.5. □

Proposition 10.3.8. *For any $f \in \overline{\mathscr{R}(\Lambda)}$,*

$$1_{\mathrm{loc}}(f) = 1_{\mathrm{loc}}(f^{\perp\perp}) \text{ and } f^{\perp\perp} = \mathcal{K}ern(1_{\mathrm{loc}}(f)).$$

Proof. $1_{\mathrm{loc}}(f^{\perp\perp}) \subseteq 1_{\mathrm{loc}}(f)$ since $f^{\perp\perp} \supseteq \langle f\rangle$. Now, let K be the $\mathcal{K}$-kernel containing f such that $1_{\mathrm{loc}}(K) = 1_{\mathrm{loc}}(f)$. K is a polar by Lemma 10.1.4(iv). By Proposition 10.1.7, $K \supseteq f^{\perp\perp}$, and so

$$1_{\mathrm{loc}}(f) = 1_{\mathrm{loc}}(K) \subseteq 1_{\mathrm{loc}}(f^{\perp\perp}) \subseteq 1_{\mathrm{loc}}(f).$$

But $f^{\perp\perp}$ is a polar and thus a $\mathcal{K}$-kernel, implying $f^{\perp\perp} = \mathcal{K}ern(1_{\mathrm{loc}}(f^{\perp\perp})) = \mathcal{K}ern(1_{\mathrm{loc}}(f))$. □

Theorem 10.3.9. *There is a* $1:1$ *correspondence*

$$\mathcal{P}lr(\overline{\mathscr{R}(\Lambda)}) \leftrightarrow 1^{\nu}\text{-Set}(\mathscr{R}^{(n)}) \tag{10.2}$$

between the polars of $\overline{\mathscr{R}(\Lambda)}$ *and the* 1^{ν}*-sets in* $\mathscr{R}^n$*, given by* $B \mapsto 1_{\mathrm{loc}}(B)$ *and* $Z \mapsto \mathcal{K}ern(Z)$*.*

This correspondence restricts to a correspondence between the principal polars of $\overline{\mathscr{R}(\Lambda)}$ *and the principal* 1^{ν}*-sets in* $\mathscr{R}^n$*.*

Proof. By Lemma 10.1.4(iv), B is a polar iff B is a $\mathcal{K}$-kernel; thus, $\mathcal{K}ern(1_{\mathrm{loc}}(B)) = B$. If $B = f^{\perp\perp}$ for some $f \in \overline{\mathscr{R}(\Lambda)}$, then $1_{\mathrm{loc}}(B) = 1_{\mathrm{loc}}(f)$ by Proposition 10.3.8, yielding $\mathcal{K}ern(1_{\mathrm{loc}}(f^{\perp\perp})) = \mathcal{K}ern(1_{\mathrm{loc}}(f)) = f^{\perp\perp}$. The restriction to principal polars follows from Lemma 10.1.10. □

Corollary 10.3.10. *Let*

$$\mathcal{B} = \{B \cap \mathscr{R}(\Lambda) : B \in \mathcal{P}lr(\overline{\mathscr{R}(\Lambda)})\}.$$

There is a 1:1 correspondence

$$\mathcal{B} \leftrightarrow 1^{\nu}\text{-Set}(\mathscr{R}^{(n)}) \tag{10.3}$$

given by $K \mapsto 1_{\mathrm{loc}}(K)$ *and* $Z \mapsto \mathcal{K}ern(Z) \cap \mathscr{R}(\Lambda) = \mathcal{K}ern_{\mathscr{R}(\Lambda)}(Z)$*, which restricts to a correspondence*

$$\text{Principal } \nu\text{-kernels of } \mathcal{B} \leftrightarrow \mathrm{P}1^{\nu}\text{-Set}\,. \tag{10.4}$$

Furthermore, $\mathcal{K}ern_{\mathscr{R}(\Lambda)}(1_{\mathrm{loc}}(K)) = \overline{K} \cap \mathscr{R}(\Lambda)$ *for any* ν*-kernel* K *of* $\mathscr{R}(\Lambda)$*.*

Proof. Since $(\mathcal{P}lr(\mathscr{R}(\Lambda)), \cdot, \cap)$ is a lattice, the first assertion follows by applying Proposition 10.3.2 to the correspondence in Theorem 10.4. The second assertion is by Theorem 10.3.9 and (10.3), and the last by Proposition 10.3.7 and Remark 8.2.7. □

Here is an analog to Proposition 5.4.3. For $X \subset \mathscr{R}^n$ we define the **restriction map** $\phi_X : \mathscr{R}(\Lambda) \to \mathscr{R}(X)$ by $f \mapsto f|_X$.

Proposition 10.3.11. *For any* 1^{ν}*-set* $X = 1_{\mathrm{loc}}(\mathcal{S}) \subseteq \mathscr{R}^n$ *with* $\mathcal{S} \subseteq \mathscr{R}(\Lambda)$*,* ϕ_X *is a homomorphism and*

$$\mathscr{R}(\Lambda)/K_{\mathcal{S}} \cong \mathscr{R}[X]. \tag{10.5}$$

where $K_{\mathcal{S}} = \mathcal{S}^{\perp\perp} \cap \mathscr{R}(\Lambda)$ *with* $\mathcal{S}^{\perp\perp}$ *is taken in the completion* $\overline{\mathscr{R}(\Lambda)}$ *of* $\mathscr{R}(\Lambda)$ *in* $Fun(\mathscr{R}^{(n)}, \mathscr{R})$*.*

Proof. Corollary 10.3.10 implies that

$$\mathcal{K}ern(\phi_X) = \{g \in \mathscr{R}(\Lambda) : g|_X = 1\} = \{g \in \mathscr{R}(\Lambda) : g \in \mathcal{S}^{\perp\perp}\} = K_{\mathcal{S}}.$$

We conclude with the isomorphism theorems 4.1.12. □

10.3.1. *The restriction of the coordinate* ν*-semifield*†*.*

Corollary 10.3.12. *For any (principal)* 1^{ν}*-set* $X = 1_{\mathrm{loc}}(\langle f\rangle) \subseteq \mathscr{R}^n$*, the restriction map of* ϕ_X *to* $\mathscr{R}$ *satisfies*

$$\langle\mathscr{R}\rangle/(\langle f\rangle \cap \langle\mathscr{R}\rangle) \cong \mathscr{R}[X]. \tag{10.6}$$

Proof. By Proposition 8.4.4, $\mathcal{K}ern(\phi_X) = \{g \in \langle\mathscr{R}\rangle : g|_X = 1\} = \{g \in \langle\mathscr{R}\rangle : g \in \langle f\rangle\} = \langle f\rangle \cap \langle\mathscr{R}\rangle$, so we conclude with the isomorphism theorems 4.1.12. □

11. Dimension

Throughout, F denotes a divisible, ν-archimedean ν-semifield†. Having the HS ν-kernels at our disposal, our next major goal is to use them to define dimension, which will be seen to be uniquely defined. We return to the notation of Construction 9.4.1, as indicated after Remark 9.4.2.

Let $\langle f\rangle \subseteq \langle F\rangle$ be a principal ν-kernel and let $\langle f\rangle = \bigcap_{i=1}^{s} K_i$, where

$$K_i = (L_i \cdot R_i) \cap \langle F\rangle = (L_i \cap \langle F\rangle) \cdot (R_i \cap \langle F\rangle) = L_i' \cdot R_i'$$

is its (full) HO-decomposition; i.e., for each $1 \le i \le s$, $R_i \in \mathcal{P}(F)$ is a region ν-kernel and $L_i \in \mathcal{P}(F)$ is either an HS-kernel or bounded from below (in which case $L_i' = \langle F\rangle$). Then by Corollary 4.1.11, we have the subdirect decomposition

$$\langle F\rangle/\langle f\rangle \hookrightarrow \prod_{i=1}^{t} \langle F\rangle/K_i = \prod_{i=1}^{t} (\langle F\rangle/L_i' \cdot R_i')$$

where $t \le s$ is the number of ν-kernels K_i for which $L_i' \neq \langle F\rangle$ (for otherwise $\langle F\rangle/K_i = \{1\}$ and can be omitted from the subdirect product).

Example 11.0.13. Consider the principal ν-kernel $\langle\lambda_1\rangle \in \mathcal{P}(F(\lambda_1,\lambda_2))$. For $\alpha \in F$ such that $\alpha > 1$, we have the following infinite strictly descending chain of principal ν-kernels

$$\langle\lambda_1\rangle \supset \langle|\lambda_1| + |\lambda_2 + 1|\rangle \supset \langle|\lambda_1| + |\alpha^{-1}\lambda_2 + 1|\rangle \supset \langle|\lambda_1| + |\alpha^{-2}\lambda_2 + 1|\rangle \supset \dots$$
$$\supset \langle|\lambda_1| + |\alpha^{-k}\lambda_2 + 1|\rangle \supset \dots$$

and the strictly ascending chain of 1^ν-sets corresponding to it.

$$1\text{-loc}(\lambda_1) \subset 1\text{-loc}(|\lambda_1| + |\lambda_2 + 1|) \subset \cdots \subset 1\text{-loc}(|\lambda_1| + |\alpha^{-k}\lambda_2 + 1|) \subset \cdots =$$
$$1\text{-loc}(\lambda_1) \subset 1\text{-loc}(\lambda_1) \cap 1\text{-loc}(\lambda_2 + 1) \subset \cdots \subset 1\text{-loc}(\lambda_1) \cap 1\text{-loc}(\alpha^{-k}\lambda_2 + 1) \subset \dots.$$

Example 11.0.14. Again, consider the principal ν-kernel $\langle x\rangle \in \mathcal{P}(F(x,y))$. Then

$$\langle x\rangle = \langle|x| + (|y+1| \wedge |\frac{1}{y} + 1|)\rangle = \langle(|x| + |y+1|) \wedge (|x| + |\frac{1}{y} + 1|)\rangle = \langle|x| + |y+1|\rangle \cap \langle|x| + |\frac{1}{y} + 1|\rangle.$$

So, we have the nontrivial decomposition of 1-loc(x) as 1-loc$(|x|+|y+1|)\cup$1-loc$(|x|+|\frac{1}{y}+1|)$ (note that 1-loc$(|x|+|y+1|) =$ 1-loc$(x)\cap$1-loc$(y+1)$, and furthermore 1-loc$(|x|+|\frac{1}{y}+1|) =$ 1-loc$(x)\cap$1-loc$(\frac{1}{y}+1)$). In a similar way, using complementary order ν-kernels, one can show that every principal ν-kernel can be nontrivially decomposed to a pair of principal ν-kernels.

11.1. **Irreducible ν-kernels.**

Examples 11.0.13 and 11.0.14 demonstrate that the lattice of principal ν-kernels $\mathcal{P}(\overline{F(\Lambda)})$ (resp. $\mathcal{P}(\langle F\rangle)$) is too rich to define reducibility or finite dimension. (See [2] for a discussion of infinite dimension.) Moreover, these examples suggest that this richness is caused by order ν-kernels. This motivates us to consider Θ-reducibility for a suitable sublattice of ν-kernels $\Theta \subset \mathcal{P}(\overline{F(\Lambda)})$ (resp. $\Theta \subset \mathcal{P}(\langle F\rangle)$).

There are various families of ν-kernels that could be utilized to define the notions of reducibility, dimensionality, and so forth. We take Θ to be the sublattice generated by HP-kernels, because of its connection to the (local) dimension of the linear spaces (in logarithmic scale) defined by the 1^ν-set corresponding to a ν-kernel. Namely, HP-kernels, and more generally HS-kernels, define affine subspaces of $F^{(n)}$ (see [29, §9.2]). We work with Definition 9.1.7.

Definition 11.1.1. A ν-kernel $\langle f\rangle \in \Omega(\overline{F(\Lambda)})$ is **reducible** if there are $\langle g\rangle, \langle h\rangle \in \Omega(\overline{F(\Lambda)})$ for which $\langle g\rangle, \langle h\rangle \not\subseteq \langle f\rangle$ but $\langle g\rangle \cap \langle h\rangle \subseteq \langle f\rangle$.

Lemma 11.1.2. *$\langle f\rangle$ is reducible iff $\langle f\rangle = \langle g\rangle \cap \langle h\rangle$ where $\langle f\rangle \neq \langle g\rangle$ and $\langle f\rangle \neq \langle h\rangle$.*

Proof. Assume $\langle f\rangle$ admits the stated condition. If $\langle f\rangle \supseteq \langle g\rangle \cap \langle h\rangle$, then $\langle f\rangle = \langle f\rangle \cdot \langle f\rangle = (\langle g\rangle \cdot \langle f\rangle) \cap (\langle h\rangle \cdot \langle f\rangle)$. Thus $\langle f\rangle = \langle g\rangle \cdot \langle f\rangle$ or $\langle f\rangle = \langle h\rangle \cdot \langle f\rangle$, implying $\langle f\rangle \supseteq \langle g\rangle$ or $\langle f\rangle \supseteq \langle g\rangle$. The converse is obvious. □

Lemma 11.1.3. *Let $\langle f\rangle$ be an HP-kernel. Then for any HP-kernels $\langle g\rangle$ and $\langle h\rangle$ such that $\langle f\rangle = \langle g\rangle \cap \langle h\rangle$ either $\langle f\rangle = \langle g\rangle$ or $\langle f\rangle = \langle h\rangle$. In other words, every HP-kernel is irreducible.*

Proof. If $\langle f\rangle = \langle g\rangle \cap \langle h\rangle$ then $\langle f\rangle \subseteq \langle g\rangle$ thus $f \in \langle g\rangle$. As both f and g are $\mathscr{L}$-monomials (up to equivalence), Lemma 9.2.1 yields $\langle f\rangle = \langle g\rangle$, which in turn, by Lemma 11.1.2, implies that $\langle f\rangle$ is irreducible. □

Corollary 11.1.4. *Any HS-kernel $\langle f\rangle$ is irreducible.*

Proof. If $\langle f\rangle = \langle g\rangle \cap \langle h\rangle$ for HP-kernels $\langle g\rangle$ and $\langle h\rangle$, then $\langle f\rangle \subseteq \langle g\rangle$. But $\langle f\rangle$ is a product $\langle f_1\rangle \cdots \langle f_t\rangle$ of finitely many HP-kernels. For each $1 \leq j \leq t$, $\langle f_j\rangle \subseteq \langle g\rangle$ yielding $\langle g\rangle = \langle f_j\rangle$ by Lemma 9.2.1, and so $\langle f\rangle = \langle g\rangle$. □

Corollary 11.1.5. *If $\langle f\rangle = \langle g\rangle \cap \langle h\rangle$ for HS-kernels $\langle g\rangle$ and $\langle h\rangle$, then either $\langle f\rangle = \langle g\rangle$ or $\langle f\rangle = \langle h\rangle$.*

Proof. Otherwise, since $\langle g\rangle$ and $\langle h\rangle$ are finite products of HP-kernels, there are HP-kernel $\langle g'\rangle \subseteq \langle g\rangle$ and $\langle h'\rangle \subseteq \langle h\rangle$ such that $\langle g'\rangle \not\subseteq \langle f\rangle$ and $\langle h'\rangle \not\subseteq \langle f\rangle$. But $\langle g'\rangle \cap \langle h'\rangle \subseteq \langle g\rangle \cap \langle h\rangle = \langle f\rangle$, implying $\langle f\rangle$ is reducible, contradicting Corollary 11.1.4. □

Proposition 11.1.6. *The irreducible ν-kernels in the lattice generated by HP-kernels are precisely the HS-kernels.*

Proof. This follows from Corollary 11.1.5, since all proper intersections in the lattice generated by HP-kernels are reducible. (Note that HP-kernels are also HS-kernels.) □

Definition 11.1.7. The **hyperspace spectrum** of $\overline{F(\Lambda)}$, denoted HSpec$(\overline{F(\Lambda)})$, is the family of irreducible ν-kernels in $\Omega(\overline{F(\Lambda)})$.

Corollary 11.1.8. HSpec$(\overline{F(\Lambda)})$ *is the family of HS-kernels in $\Omega(\overline{F(\Lambda)})$, which is precisely the family of HS-fractions of $\overline{F(\Lambda)}$.*

Definition 11.1.9. A chain $P_0 \subset P_1 \subset \cdots \subset P_t$ in HSpec$(\overline{F(\Lambda)})$ of HS-kernels of $\overline{F(\Lambda)}$ is said to have **length** t. An HS-kernel P has **height** t (denoted hgt$(P) = t$) if there is a chain of length t in HSpec$(\overline{F(\Lambda)})$ terminating at P, but no chain of length $t+1$ terminates at P.

Remark 11.1.10. Let L be a ν-kernel in $\mathcal{P}(\overline{F(\Lambda)})$. Consider the canonical homomorphism $\phi_L : \overline{F(\Lambda)} \to \overline{F(\Lambda)}/L$. Since the image of a principal ν-kernel is generated by the image of any of its generators, $\phi_L(\langle f\rangle) = \langle \phi_L(f)\rangle$ for any HP-kernel $\langle f\rangle$. Choosing f to be an $\mathscr{L}$-monomial, $\langle \phi_L(f)\rangle$ is a nontrivial HP-kernel in $\overline{F(\Lambda)}/L$ if and only if $\phi_L(f) \notin F$. Thus, the set of HP-kernels of $\overline{F(\Lambda)}$ mapped to HP-kernels of $\overline{F(\Lambda)}/L$ is

$$\left\{\langle g\rangle : \langle g\rangle \cdot \langle F\rangle \supseteq \phi_L^{-1}\left(\langle F\rangle\right) = L \cdot \langle F\rangle\right\}. \tag{11.1}$$

As ϕ_L is an F-homomorphism, it respects $\vee, \wedge$ and $|\cdot|$, and thus $\phi_L((\Omega(\overline{F(\Lambda)}), \cap, \cdot)) = (\Omega(\overline{F(\Lambda)}/L), \cap, \cdot)$. In fact Theorem 4.1.10 yields a correspondence identifying $\mathrm{HSpec}(\overline{F(\Lambda)}/L)$ with the subset of $\mathrm{HSpec}(\overline{F(\Lambda)})$ which consists of all HS-kernels P of $\overline{F(\Lambda)}$ such that $P \cdot \langle F\rangle \supseteq L \cdot \langle F\rangle$.

Lemma 11.1.11. *The above correspondence extends to a correspondence identifying $\Omega(\overline{F(\Lambda)}/L)$ with the subset (11.1) of $\Omega(\overline{F(\Lambda)})$. Under this correspondence, the maximal HS-kernels of $\overline{F(\Lambda)}/L$ correspond to maximal HS-kernels of $\overline{F(\Lambda)}$, and reducible ν-kernels of $\overline{F(\Lambda)}/L$ correspond to reducible ν-kernels of $\overline{F(\Lambda)}$.*

Proof. The latter assertion is obvious since $\wedge$ is preserved under homomorphisms. For the first assertion, $(\overline{F(\Lambda)}/L)/(P/L) \cong \overline{F(\Lambda)}/P$ by Theorem 4.1.12, so simplicity of the quotients is preserved. Hence, so is maximality of P/L and P. □

Definition 11.1.12. The Hyperdimension of $\overline{F(\Lambda)}$, written $\mathrm{Hdim}\,\overline{F(\Lambda)}$ (if it exists), is the maximal height of the HS-kernels in $\overline{F(\Lambda)}$.

11.2. **Decompositions.**

Let us garner some information about reducible ν-kernels, from rational functions. Suppose $f \in \overline{F(\Lambda)}$. We write $f = \sum_{i=1}^k f_i$ where each f_i is of the form $g_i {h_i}^*$ with $g_i, h_i \in F[\Lambda]$ and g_i a monomial. (Thus $f_i(\mathbf{a}) = \frac{g_i(\mathbf{a})}{h_i(\mathbf{a})}$ when $h_i(\mathbf{a})$ is tangible.) We also assume that this sum is **irredundant** in the sense that we cannot remove any of the summands and still get f. If each time the value 1 is attained by one of the terms f_i in this expansion and all other terms attain values ≤ 1, then $\tilde{f} = \bigwedge_{i=1}^k |f_i|$ defines the same 1^ν-set as f. Moreover, if $f \in \langle F\rangle$ then $\tilde{f} \wedge |\alpha| \in \langle F\rangle$, for $\alpha \in F \setminus \{1\}$ is also a generator of $\langle f\rangle$. The reason we take $\tilde{f} \wedge |\alpha|$ is that we have no guarantee that each of the f_i's in the above expansion is bounded.

We can generalize this idea as follows:

We call $f \in \overline{F(\Lambda)}$ **reducible** if we can write $f = \sum_{i=1}^k f_i$ as above, such that for every $1 \leq i \leq k$ the following condition holds:

$$f_i(\mathbf{a}) \cong_\nu 1 \Rightarrow \quad f_j(\mathbf{a}) \leq_\nu 1, \ \forall j \neq i.$$

Definition 11.2.1. Let $f \in \overline{F(\Lambda)}$. A Θ-**decomposition** of f is an expression of the form

$$|f| = |u| \wedge |v| \tag{11.2}$$

with u, v Θ-elements in $\overline{F(\Lambda)}$.

The decomposition (11.2) is said to be trivial if $f \sim_{\langle F\rangle} u$ or $f \sim_{\langle F\rangle} v$ (equivalently $|f| \sim_{\langle F\rangle} |u|$ or $|f| \sim_{\langle F\rangle} |v|$). Otherwise, the decomposition is said to be **nontrivial**.

Lemma 11.2.2. *Suppose $f \in \overline{F(\Lambda)}$ is a Θ-element. Then $\langle f\rangle$ is reducible if and only if there exists some generator f' of $\langle f\rangle$ that has a nontrivial Θ-decomposition.*

Proof. If $\langle f\rangle$ is reducible , then there exist ν-kernels $\langle u\rangle$ and $\langle v\rangle$ in Θ such that $\langle f\rangle = \langle u\rangle \cap \langle v\rangle$ where $\langle f\rangle \neq \langle u\rangle$ and $\langle f\rangle \neq \langle v\rangle$. Since $\langle u\rangle \cap \langle v\rangle = \langle |u| \wedge |v|\rangle$ we have the nontrivial Θ-decomposition $f' = |u| \wedge |v|$ (which is a generator of $\langle f\rangle$).

Conversely, assume that $f' = |u| \wedge |v|$ is a nontrivial Θ- decomposition for some $f' \sim_{\langle F\rangle} f$. Then $\langle f\rangle = \langle f'\rangle = \langle |u| \wedge |v|\rangle = \langle u\rangle \cap \langle v\rangle$. Since the decomposition $f' = |u| \wedge |v|$ is nontrivial, we have that $u \not\sim_{\langle F\rangle} f'$ and $v \not\sim_{\langle F\rangle} f'$, and thus $\langle |u|\rangle = \langle u\rangle \neq \langle f'\rangle = \langle f\rangle$. Similarly, $\langle v\rangle \neq \langle f\rangle$. Thus, by definition, $\langle f\rangle$ is reducible. □

We can equivalently rephrase Lemma 11.2.2 as follows:

Remark 11.2.3. f is reducible if and only if some $f' \sim_{\langle F\rangle} f$ has a nontrivial Θ-decomposition.

A question immediately arising from Definition 11.2.1 and Lemma 11.2.2 is:

If $f \in \overline{F(\Lambda)}$ has a nontrivial Θ-decomposition and $g \sim_{\langle F\rangle} f$, does g also have a nontrivial Θ-decomposition? If so, how is this pair of decompositions related?

In the next few paragraphs we provide an answer to both of these questions, for $\Theta = \mathcal{P}(\langle \mathscr{R}\rangle)$.

Remark 11.2.4. $\sum_{i=1}^k s_i(a_i \wedge b_i)^{d(i)} = \left(\sum_{i=1}^k s_i a_i^{d(i)}\right) \wedge \left(\sum_{i=1}^k s_i b_i^{d(i)}\right)$, $\forall s_1, ..., s_k, a_1, ..., a_k, b_1, ..., b_k \in \overline{F(\Lambda)}$, and $d(i) \in \mathbb{N}_{\geq 0}$.

Remark 11.2.5. If $h_1, ..., h_k \in \overline{F(\Lambda)}$ such that each $h_i \geq_\nu 1$, then $\sum_{i=1}^k s_i h_i \geq_\nu 1$ for every $s_1, ..., s_k \in \overline{F(\Lambda)}$ such that $\sum_{i=1}^k s_i \cong_\nu 1$.

Theorem 11.2.6. *(For $\Theta = \mathcal{P}(\langle \mathscr{R}\rangle)$.) If $\langle f\rangle$ is a (principal) reducible ν-kernel, then there exist Θ-elements $g, h \in \overline{F(\Lambda)}$ such that $|f| = |g| \wedge |h|$ and $|f| \not\sim_{\langle F\rangle} |g|, |h|$.*

Proof. If $\langle f\rangle$ is a principal reducible ν-kernel, then there exists $f' \sim_{\langle F\rangle} f$ such that $f' = |u| \wedge |v| = \min(|u|, |v|)$ for Θ-elements $u, v \in \langle \mathscr{R}\rangle$ with $f' \not\sim_{\langle F\rangle} |u|, |v|$. Then $|f| \in \langle f'\rangle$ since f' is a generator of $\langle f\rangle$, so there exist $s_1, ..., s_k \in \overline{F(\Lambda)}$ such that $\sum_{i=1}^k s_i = 1$ and $|f| = \sum_{i=1}^k s_i (f')^{d(i)}$ with $d(i) \in \mathbb{N}_{\geq 0}$. ($d(i) \geq 0$ since $|f| \geq_\nu 1$.) Thus

$$f = \sum_{i=1}^k s_i (|u| \wedge |v|)^{d(i)} = \sum_{i=1}^k s_i (\min(|u|, |v|))^{d(i)} = \min\left(\sum_{i=1}^k s_i |u|^{d(i)}, \sum_{i=1}^k s_i |v|^{d(i)}\right) = |g| \wedge |h|$$

where $g = |g| = \sum_{i=1}^k s_i |u|^{d(i)}, h = |h| = \sum_{i=1}^k s_i |v|^{d(i)}$.

Now $\langle |f|\rangle \subseteq \langle |g|\rangle \subseteq \langle |u|\rangle$ and $\langle |f|\rangle \subseteq \langle |h|\rangle \subseteq \langle |v|\rangle$, implying $1_{\mathrm{loc}}(f) \supseteq 1_{\mathrm{loc}}(g) \supseteq 1_{\mathrm{loc}}(u)$ and $1_{\mathrm{loc}}(f) \supseteq 1_{\mathrm{loc}}(h) \supseteq 1_{\mathrm{loc}}(v)$.

We claim that $|g|$ and $|h|$ generate $\langle |u|\rangle$ and $\langle |v|\rangle$, respectively. Indeed, $1_{\mathrm{loc}}(f') = 1_{\mathrm{loc}}(|f|)$, since $f' \sim_{\langle F\rangle} |f|$ and thus for any $\mathbf{a} \in F^{(n)}$, $f'(\mathbf{a}) = 1 \Leftrightarrow |f|(\mathbf{a}) = 1$. Let $s_j (f')^{d(j)}$ be a dominant term of $|f|$ at $\mathbf{a}$, i.e.,

$$|f| \cong_\nu \sum_{i=1}^k s_i(\mathbf{a})(f'(\mathbf{a}))^{d(i)} \cong_\nu s_j(\mathbf{a})(f'(\mathbf{a}))^{d(j)}.$$

Then $f(\mathbf{a}) \cong_\nu 1 \Leftrightarrow s_j(\mathbf{a})(f'(\mathbf{a}))^{d(j)} \cong_\nu 1$. If $f'(\mathbf{a}) \cong_\nu 1$, then $(f'(\mathbf{a}))^{d(j)} \cong_\nu 1$, so $s_j(\mathbf{a}) \cong_\nu 1$. Now, for $\mathbf{a} \in 1_{\mathrm{loc}}(g)$. Then we have

$$g(\mathbf{a}) \cong_\nu \sum_{i=1}^k s_i |u|^{d(i)} \cong_\nu 1.$$

Let $s_t |u|^{d(t)}$ be a dominant term of g at $\mathbf{a}$. If $s_t(\mathbf{a}) \cong_\nu 1$ then $|u|^{d(t)} \cong_\nu 1$ and thus $u = 1$, and $\mathbf{a} \in 1_{\mathrm{loc}}(u)$. Otherwise $s_t(\mathbf{a}) <_\nu 1$ (since $\sum_{i=1}^k s_i \cong_\nu 1$) and so, by the above, $s_t (f')^{d(t)}$ is not a dominant term of $|f|$ at $\mathbf{a}$. Thus, for every index j of a dominant term of $|f|$ at $\mathbf{a}$, we have $j \neq t$ and

$$|u(\mathbf{a})|^{d(j)} \cong_\nu s_j(\mathbf{a})|u(\mathbf{a})|^{d(j)} <_\nu s_t(\mathbf{a})|u(\mathbf{a})|^{d(t)} \cong_\nu g(\mathbf{a}) \cong_\nu 1.$$

$$s_j(\mathbf{a})(f'(\mathbf{a}))^{d(j)} \cong_\nu s_j(\mathbf{a})(|u|(\mathbf{a}) \wedge |v|(\mathbf{a}))^{d(j)} \leq s_j(\mathbf{a})|u(\mathbf{a})|^{d(j)} <_\nu 1. \tag{11.3}$$

On the other hand, $f'(\mathbf{a}) \cong_\nu 1$ since $1_{\mathrm{loc}}(f) \supseteq 1_{\mathrm{loc}}(g)$, implying $s_j(\mathbf{a})(f'(\mathbf{a}))^{d(j)} \cong_\nu 1$, contradicting (11.3). Hence, $1_{\mathrm{loc}}(g) \subseteq 1_{\mathrm{loc}}(u)$, yielding $1_{\mathrm{loc}}(g) = 1_{\mathrm{loc}}(u)$, which implies that g is a generator of $\langle |u|\rangle = \langle u\rangle$. The proofs for h and $|v|$ are analogous.

Consequently, $g \sim_{\langle F\rangle} |g| \sim_{\langle F\rangle} |u|$ and $h \sim_{\langle F\rangle} |h| \sim_{\langle F\rangle} |v|$. Since $|f| \sim_{\langle F\rangle} f' \not\sim_{\langle F\rangle} |u|, |v|$ we conclude that $|f| \not\sim_{\langle F\rangle} |g|, |h|$. □

Corollary 11.2.7. *For $f \in \langle F\rangle$, if $|f| = \bigwedge_{i=1}^s |f_i|$ for $f_i \in \langle F\rangle$, then for any $g \sim_{\langle F\rangle} f$, we have $|g| = \bigwedge_{i=1}^s |g_i|$, with $g_i \sim_{\langle F\rangle} f_i$ for $i = 1, ..., s$.*

Proof. Iterate Theorem 11.2.6. □

Corollary 11.2.8. *If $\langle f\rangle$ is a ν-kernel in Θ, then $\langle f\rangle$ has a nontrivial decomposition $\langle f\rangle = \langle g\rangle \cap \langle h\rangle$ if and only if $|f|$ has a nontrivial decomposition $|f| = |g'| \wedge |h'|$ with $|g'| \sim_{\langle F\rangle} g$ and $|h'| \sim_{\langle F\rangle} h$.*

Proof. If $|f| = |g'| \wedge |h'|$ then, since $|g'| \sim_{\langle F\rangle} g$ and $|h'| \sim_{\langle F\rangle} h$ we have

$$\langle f\rangle = \langle |f|\rangle = \langle |g'| \wedge |h'|\rangle = \langle |g'|\rangle \cap \langle |h'|\rangle = \langle g\rangle \cap \langle h\rangle.$$

The converse is seen as in the proof of Theorem 11.2.6. □

Corollary 11.2.8 provides a Θ-decomposition of $|f|$, for every generator f of a reducible ν-kernel in Θ.

Remark 11.2.9. By [29, Corollary 4.1.25],

$$\langle f\rangle \cap \langle g\rangle = \langle (f + f^*) \wedge (g + g^*)\rangle = \langle |f| \wedge |g|\rangle.$$

But, in fact, $\langle f\rangle \cap \langle g\rangle = \langle f'\rangle \cap \langle g'\rangle$ for any $g' \sim_{\langle F\rangle} g$ and $h' \sim_{\langle F\rangle} h$, so we could take $|g'| \wedge |f'|$ instead of $|g| \wedge |f|$ on the righthand side of the equality, e.g., $\langle |f^k| \wedge |g^m|\rangle$ for any $m, k \in \mathbb{Z} \setminus \{0\}$.

Definition 11.2.10. Let $\mathbb{S}$ be a semifield and let $a, b \in \mathbb{S}$. We say that a and b are $\langle F\rangle$-**comparable** if there exist some $a' \sim_{\langle F\rangle} a$ and $b' \sim_{\langle F\rangle} b$ such that $|a'| \leq |b'|$ or $|b'| \leq |a'|$.

Since $|g| \wedge |h| = \min(|g|, |h|)$ we can utilize Remark 11.2.9 to get the following observation:

Proposition 11.2.11. *A Θ-decomposition $f \sim_{\langle F\rangle} |g| \wedge |h| \in \overline{F(\Lambda)}$ is nontrivial if and only if the Θ-elements g and h are not $\langle F\rangle$-comparable.*

Proof. If g and h are $\langle F\rangle$-comparable, then there exist some $g' \sim_{\langle F\rangle} g$ and $h' \sim_{\langle F\rangle} h$ such that $|g'| \geq_\nu |h'|$ or $|h'| \geq_\nu |g'|$. Without loss of generality, assume that $|g'| \geq_\nu |h'|$. Then $\langle |g| \wedge |h|\rangle = \langle |g|\rangle \cap \langle |h|\rangle = \langle |g'|\rangle \cap \langle |h'|\rangle = \langle |g'| \wedge |h'|\rangle = \langle \min(|g'|, |h')|\rangle = \langle |g'|\rangle = \langle g'\rangle = \langle g\rangle$. Thus $\langle f\rangle = \langle g\rangle$ so $f \sim_{\langle F\rangle} g$ yielding that the decomposition is trivial.

Conversely, if g and h are not $\langle F\rangle$-comparable then we claim that $f \not\sim_{\langle F\rangle} g$ and $f \not\sim_{\langle F\rangle} h$. We must show that $\langle h\rangle \not\subseteq \langle g\rangle$ and $\langle g\rangle \not\subseteq \langle h\rangle$ respectively. So assume that $\langle h\rangle \supseteq \langle g\rangle$. In view of Lemma 9.2.1 and [29, Proposition 4.1.13], there exists some $f' \sim f$ such that $f' = |h|^k \wedge g$. Note that $|h|^k \geq 1$ and $g \geq 1$ so $f' = |h|^k \wedge g \geq 1$ and thus $|f'| = f'$. Finally,

$$|f'| = |h|^k \wedge g \Leftrightarrow |f'| \leq |h|^k \Leftrightarrow |f'| \in \langle |h|\rangle \Leftrightarrow f' \in \langle h\rangle \Leftrightarrow f \in \langle h\rangle.$$

□

11.3. **Convex dependence.**

Definition 11.3.1. An HS-fraction f of $\overline{F(\Lambda)}$ is F**-convexly dependent** on a set A of HS-fractions if

$$f \in \langle \{g : g \in A\}\rangle \cdot \langle F\rangle; \tag{11.4}$$

otherwise f is said to be F**-convexly independent** of A. The set A is said to be F**-convexly independent** if f is F-convexly independent of $A \setminus \{f\}$, for every $f \in A$. If $\{a_1, ..., a_n\}$ is F-convexly dependent, then we also say that $a_1, ..., a_n$ are F-convexly dependent.

Note that under the assumption that $g \in \langle F\rangle \setminus \{1\}$ for some $g \in A$, the condition in (11.4) simplifies to $f \in \langle \{g : g \in A\}\rangle$.

Remark 11.3.2. By definition, an HS-fraction f is F-convexly dependent on HS-fractions $\{g_1, ..., g_t\}$ if and only if

$$\langle |f|\rangle = \langle f\rangle \subseteq \langle g_1, ..., g_t\rangle \cdot \langle F\rangle = \Big\langle \sum_{i=1}^{t} |g_i| \Big\rangle \cdot \langle F\rangle = \Big\langle \sum_{i=1}^{t} |g_i| + |\alpha| \Big\rangle,$$

for any element α of F for which $\alpha^\nu \neq 1^\nu$.

Example 11.3.3. For any $\alpha \in F$ and any $f \in \overline{F(\Lambda)}$,

$$|\alpha f| \leq |f|^2 + |\alpha|^2 = (|f| + |\alpha|)^2.$$

Thus $\alpha f \in \big\langle (|f| + |\alpha|)^2 \big\rangle = \langle |f| + |\alpha|\rangle = \langle f\rangle \cdot \langle F\rangle$. In particular, if f is an HS-fraction, then αf is F-convexly dependent on f.

As a consequence of Lemma 9.2.1, if two $\mathscr{L}$-monomials f, g, satisfy $g \in \langle f\rangle$, then $\langle g\rangle = \langle f\rangle$. In other words, either $\langle g\rangle = \langle f\rangle$ or $\langle g\rangle \not\subseteq \langle f\rangle$ and $\langle f\rangle \not\subseteq \langle g\rangle$. This motivates us to restrict the convex dependence relation to the set of $\mathscr{L}$-monomials. This will be justified later by showing that for each F-convexly independent subset of HS-fractions of order t in $\overline{F(\Lambda)}$, there exists an F-convexly independent subset of $\mathscr{L}$-monomials having order $\geq t$ in $\overline{F(\Lambda)}$. Let us see that convex-dependence is an abstract dependence relation.

Proposition 11.3.4. *Let $A, A_1 \subset \overline{F(\Lambda)}$ be sets of HS-fractions, and let f be an HS-fraction.*

1. *If $f \in A$, then f is F-convexly dependent on A.*
2. *If f is F-convexly dependent on A and each $a \in A$ is F-convexly dependent on A_1, then f is F-convexly dependent on A_1.*
3. *If f is F-convexly dependent on A, then f is F-convexly dependent on A_0 for some finite subset A_0 of A.*

Proof. (1) $f \in \langle A\rangle \subseteq \langle A\rangle \cdot \langle F\rangle$.

(2) $\langle A\rangle \subseteq \langle A_1\rangle \cdot \langle F\rangle$ since a is convexly dependent on A_1 for each $a \in A$. If f is F-convexly dependent on A, then $f \in \langle A\rangle \cdot \langle F\rangle \subseteq \langle A_1\rangle \cdot \langle F\rangle$, so, f is F-convexly dependent on A_1.

(3) $a \in \langle A\rangle \cdot \langle F\rangle$, so by Proposition **??** there exist some $s_1, ..., s_k \in \overline{F(\Lambda)}$ and $g_1, ..., g_k \in G(A \cup F) \subset \langle A\rangle \cdot \langle F\rangle$, where $G(A \cup F)$ is the group generated by $A \cup F$, such that $\sum_{i=1}^k s_i = 1$ and $a = \sum_{i=1}^k s_i g_i^{d(i)}$ with $d(i) \in \mathbb{Z}$. Thus $a \in \langle g_1, ..., g_k\rangle$ and $A_0 = \{g_1, ..., g_k\}$. □

From now on, we assume that the ν-semifield† F is divisible.

Proposition 11.3.5 (Steinitz exchange axiom)**.** *Let $S = \{b_1, ..., b_t\} \subset \mathrm{HP}(\overline{F(\Lambda)})$ and let f and b be elements of $\mathrm{HP}(\overline{F(\Lambda)})$. If f is F-convexly dependent on $S \cup \{b\}$ and f is F-convexly independent of S, then b is F-convexly dependent on $S \cup \{f\}$.*

Proof. We may assume that $\alpha \in S$ for some $\alpha \in F$. Since f is F-convexly independent of S, by definition $f \notin \langle S\rangle$ this implies that $\langle S\rangle \subset \langle S\rangle \cdot \langle f\rangle$ (for otherwise $\langle f\rangle \subseteq \langle S\rangle$ yielding that f is F-convexly dependent on S). Since f is F-convexly dependent on $S \cup \{b\}$, we have that $f \in \langle S \cup \{b\}\rangle = \langle S\rangle \cdot \langle b\rangle$. In particular, we get that $b \notin \langle S\rangle \cdot \langle F\rangle$ for otherwise f would be dependent on S. Consider the quotient map $\phi : \overline{F(\Lambda)} \to \overline{F(\Lambda)}/\langle S\rangle$. Since ϕ is a semifield epimorphism and $f, b \notin \langle S\rangle \cdot \langle F\rangle = \phi^{-1}(\langle F\rangle)$, we have that $\phi(f)$ and $\phi(b)$ are not in F thus are $\mathscr{L}$-monomials in the semifield $\mathrm{Im}\,(\phi) = \overline{F(\Lambda)}/\langle S\rangle$. By the above, $\phi(f) \neq 1$ and $\phi(f) \in \phi(\langle b\rangle) = \langle \phi(b)\rangle$. Thus, $\langle \phi(f)\rangle = \langle \phi(b)\rangle$ by Lemma 9.2.1. So $\langle S\rangle \cdot \langle f\rangle = \phi^{-1}(\langle \phi(f)\rangle) = \phi^{-1}(\langle \phi(b)\rangle) = \langle S\rangle \cdot \langle b\rangle$, consequently $b \in \langle S\rangle \cdot \langle b\rangle = \langle S\rangle \cdot \langle f\rangle = \langle S \cup \{f\}\rangle$, i.e., b is F-convexly dependent on $S \cup \{f\}$. □

Definition 11.3.6. Let $A \subseteq \mathrm{HP}(\overline{F(\Lambda)})$. The **convex span** of A over F is the set

$$\mathrm{Conv}_F(A) = \{a \in \mathrm{HP}(\overline{F(\Lambda)}) : a \text{ is } F\text{-convexly dependent on } A\}. \tag{11.5}$$

For a ν-semifield† $\mathrm{K} \subseteq \overline{F(\Lambda)}$ such that $F \subseteq \mathrm{K}$. a set $A \subseteq \mathrm{HP}(\overline{F(\Lambda)})$ is said to **convexly span** K over F if

$$\mathrm{HP}(\mathrm{K}) = \mathrm{Conv}_F(A).$$

Remark 11.3.7. $\mathrm{Conv}(\{f_1, ..., f_m\}) = \langle f_1, ..., f_m\rangle \cdot \langle F\rangle$.

In view of Propositions 11.3.4 and 11.3.5, convex dependence on $\mathrm{HP}(\overline{F(\Lambda)})$ is an abstract dependence relation. Then by [33, Chapter 6], we have:

Corollary 11.3.8. *Let $V \subset \mathrm{HP}(\overline{F(\Lambda)})$. Then any set convexly spanning V contains a basis $B_V \subset V$, which is a maximal convexly independent subset of unique cardinality such that*

$$\mathrm{Conv}(B_V) = \mathrm{Conv}(V).$$

Example 11.3.9. By Lemma 8.1.17, the maximal ν-kernels in $\mathcal{P}(\overline{F(\Lambda)})$ are HS-fractions of the form $L_{(\alpha_1,\dots,\alpha_n)} = \langle \alpha_1 x_1, \dots, \alpha_n x_n\rangle$ for any $\alpha_1, ..., \alpha_n \in F$. In view of Corollary 11.3.8,

$$\overline{F(\Lambda)} = \mathrm{Conv}(\{\alpha_1 x_1, \dots, \alpha_n x_n\}),$$

i.e., $\{\alpha_1 x_1, \dots, \alpha_n x_n\}$ convexly spans $\overline{F(\Lambda)}$ over F. Now $\alpha_k x_k \notin \langle \bigcup_{j\neq k} \alpha_j x_j\rangle \cdot \langle F\rangle$, since there are no order relations between $\alpha_i x_i$ and the elements of $\{\alpha_j x_j : j \neq i\} \cup \{\alpha \ : \alpha \in F\}$. Thus, for arbitrary $\alpha_1, ..., \alpha_n \in F$, $\{\alpha_1 x_1, \dots, \alpha_n x_n\}$ is F-convexly independent, constituting a basis for $\overline{F(\Lambda)}$.

Definition 11.3.10. Let $V \subset \mathrm{HP}(\overline{F(\Lambda)})$ be a set of $\mathscr{L}$-monomials. We define the **convex dimension** of V, $\mathrm{d}_{\mathrm{conv}}(V)$, to be $|B|$ where B is a basis for V.

Example 11.3.11. $\mathrm{d}_{\mathrm{conv}}\left(\overline{F(\Lambda)}\right) = n$, by Example 11.3.9.

Remark 11.3.12. If $S \subset \mathrm{HP}(\overline{F(\Lambda)})$, then for any $f, g \in \overline{F(\Lambda)}$ such that $f, g \in \mathrm{Conv}(S)$

$$|f| + |g| \in \mathrm{Conv}(S) \ \text{ and } \ |f| \wedge |g| \in \mathrm{Conv}(S).$$

Proof. First we prove that $|f| + |g| \in \mathrm{Conv}(S)$. Since $\langle f\rangle \subseteq \langle S\rangle \cdot \langle F\rangle$ and $\langle g\rangle \subseteq \langle S\rangle \cdot \langle F\rangle$, we have $\langle f, g\rangle = \langle |f| + |g|\rangle = \langle f\rangle \cdot \langle g\rangle \subseteq \langle S\rangle \cdot \langle F\rangle$. $|f| \wedge |g| \in \mathrm{Conv}(S)$, since $\langle |f| \wedge |g|\rangle = \langle f\rangle \cap \langle g\rangle \subseteq \langle g\rangle \subseteq \langle S\rangle \cdot \langle F\rangle$. □

Remark 11.3.13. If K is an HS-kernel, then K is generated by an HS-fraction $f \in \overline{F(\Lambda)}$ of the form $f = \sum_{i=1}^t |f_i|$ where $f_1, ..., f_t$ are $\mathscr{L}$-monomials. So,

$$\mathrm{Conv}(K) = \langle F\rangle \cdot K = \langle F\rangle \cdot \langle f\rangle = \langle F\rangle \cdot \Big\langle \sum_{i=1}^t |f_i| \Big\rangle = \langle F\rangle \cdot \prod_{i=1}^t \langle f_i\rangle$$
$$= \langle F\rangle \cdot \langle f_1, ..., f_t\rangle$$

and so, $\{f_1, ..., f_t\}$ convexly spans $\langle F\rangle \cdot K$.

Remark 11.3.14. Let f be an HS-fraction. Then $f \sim_{\langle F\rangle} \sum_{i=1}^t |f_i|$ where f_i are $\mathscr{L}$-monomials. Hence f is F-convexly dependent on $\{f_1, ..., f_t\}$, since $\langle f\rangle = \prod_{i=1}^t \langle f_i\rangle = \langle \{f_1, ..., f_t\}\rangle$.

Lemma 11.3.15. *Suppose $\{b_1, ..., b_m\}$ is a set of HS-fractions, such that $b_i \sim_{\langle F\rangle} \sum_{j=1}^{t_i} |f_{i,j}|$, where $f_{i,j}$ are $\mathscr{L}$-monomials. Then b_1 is F-convexly dependent on $\{b_2, ..., b_m\}$ if and only if all of its summands $f_{1,r}$ for $1 \leq r \leq t_1$ are F-convexly dependent on $\{b_2, ..., b_m\}$.*

Proof. If b_1 is F-convexly dependent on $\{b_2, ..., b_m\}$, then

$$\prod_{j=1}^{t_1} \langle f_{1,j}\rangle = \Big\langle \sum_{j=1}^{t_1} |f_{1,j}| \Big\rangle = \langle b_1\rangle \subseteq \big\langle \{b_1, ..., b_m\}\big\rangle.$$

Hence $f_{1,r} \in \prod_{j=1}^{t_1} \langle f_{1,j}\rangle$ is F-convexly dependent on $\{b_1, ..., b_m\}$ and by Remark 11.3.14 $f_{1,r}$ is F-convexly dependent on $\{f_{i,j} : 2 \leq i \leq m;\ 1 \leq j \leq t_i\}$. Conversely, if each $f_{1,r}$ is F-convexly dependent on $\{b_2, ..., b_m\}$ for $1 \leq r \leq t_1$, then there exist some $k_1, ..., k_{t_1}$ such that $|f_{1,r}| \leq \sum_{i=2}^m \sum_{j=1}^{t_i} |f_{i,j}|^{k_r}$. Hence b_1 is F-convexly dependent on $\{b_2, ..., b_m\}$, by Corollary 4.1.21. □

Lemma 11.3.16. *Let $V = \{f_1, ..., f_m\}$ be a F-convexly independent set of HS-fractions, with $f_i \sim_{\langle F\rangle} \sum_{j=1}^{t_i} |f_{i,j}|$ for $\mathscr{L}$-monomials $f_{i,j}$. Then there exists an F-convexly independent subset*

$$S_0 \subseteq S = \{f_{i,j} : 1 \leq i \leq m; 1 \leq j \leq t_i\}$$

such that $|S_0| \geq |V|$ and $\mathrm{Conv}(S_0) = \mathrm{Conv}(V)$.

Proof. By Remark 11.3.14, f_i is dependent on the set of $\mathscr{L}$-monomials $\{f_{i,j} : 1 \leq j \leq t_i\} \subset S$ for each $1 \leq i \leq m$, implying $\mathrm{Conv}(S) = \mathrm{Conv}(V)$. By Corollary 11.3.8, S contains a maximal F-convexly independent subset S_0 such that $\mathrm{Conv}(S_0) = \mathrm{Conv}(S)$ which, by Lemma 11.3.15, we can shrink down to a base. □

Lemma 11.3.17. *The following hold for an $\mathscr{L}$-monomial f:*

(1) $\langle F\rangle \not\subseteq \langle f\rangle$.
(2) *If $\overline{F(\Lambda)}$ is not bounded, then $\langle f\rangle \not\subseteq \langle F\rangle$.*

Proof. By definition, an $\mathscr{L}$-monomial is not bounded from below. Thus $\langle f\rangle \cap F = \{1\}$, yielding $\langle F\rangle \not\subseteq \langle f\rangle$. For the second assertion, an HP-kernel is not bounded when $\overline{F(\Lambda)}$ is not bounded, so $\langle f\rangle \not\subseteq \langle F\rangle$. □

A direct consequence of Lemma 11.3.17 is:

Lemma 11.3.18. *If $\overline{F(\Lambda)}$ is not bounded, then any nontrivcial HS-kernel (i.e., $\neq \langle 1\rangle$) is F-convexly independent.*

Proof. By Lemma 11.3.17 the assertion is true for HP-kernels, and thus for HS-kernels, since every HS-kernel contains some HP-kernel. □

11.4. **Computing convex dimension.**

Having justified our restriction to $\mathscr{L}$-monomials, we move ahead with computing lengths of chains.

Remark 11.4.1. Let K be an HS-kernel of $\overline{F(\Lambda)}$. By definition there are $\mathscr{L}$-monomials $f_1, ..., f_t \in \mathrm{HP}(K)$ such that $K = \langle \sum_{i=1}^t |f_i|\rangle$. By Remark 11.3.13, $\mathrm{Conv}(K)$ is convexly spanned by $f_1, ..., f_t$. Now, since $\mathrm{Conv}(K) = \mathrm{Conv}(f_1, ..., f_t)$ and $\{f_1, ..., f_t\} \subset \mathrm{HP}(K) \subset \mathrm{HP}(\overline{F(\Lambda)})$, by Corollary 11.3.8, $\{f_1, ..., f_t\}$ contains a basis $B = \{b_1, ..., b_s\} \subset \{f_1, ..., f_t\}$ of F-convexly independent elements, where $s = \mathrm{d}_{\mathrm{conv}}(K)$, such that $\mathrm{Conv}(B) = \mathrm{Conv}(f_1, ..., f_t) = \mathrm{Conv}(K)$.

Proposition 11.4.2. *For any order ν-kernel o of $\overline{F(\Lambda)}$, if $\mathscr{L}$-monomials $h_1, ..., h_t$ are F-convexly dependent, then the images of $h_1, ..., h_t$ are F-convexly dependent (in the quotient ν-semifield† $\overline{F(\Lambda)}/o$).*

Proof. Denote by $\phi_o : \overline{F(\Lambda)} \to \overline{F(\Lambda)}/o$ the quotient F-homomorphism. Then $\phi_o(\langle F\rangle) = \langle \phi_o(F)\rangle = \langle F\rangle_{\overline{F(\Lambda)}/o}$. Now, if $h_1, ..., h_t$ are F-convexly dependent then there exist some j, say without loss of generality $j = 1$, such that $h_1 \in \langle h_2, ..., h_t\rangle \cdot \langle F\rangle$. By assumption and Proposition 4.1.18,

$$\begin{aligned}\phi_o(h_1) &\in \phi_o\left(\langle h_2, ..., h_t, \alpha\rangle\right)\\ &= \langle \phi_o(h_2), ..., \phi_o(h_t), \phi_o(\alpha)\rangle = \langle \phi_o(h_2), ..., \phi_o(h_t), \alpha\rangle\\ &= \langle \phi_o(h_2), ..., \phi_o(h_t)\rangle \cdot \langle F\rangle\end{aligned}$$

Thus $\phi_o(h_1)$ is F-convexly dependent on $\{\phi_o(h_2), ..., \phi_o(h_t)\}$. □

Conversely, we have:

Lemma 11.4.3. *For any order ν-kernel o of $\overline{F(\Lambda)}$, and any set $\{h_1, ..., h_t\}$ of $\mathscr{L}$-monomials, if $\phi_o(h_1), ..., \phi_o(h_t)$ are F-convexly dependent in the quotient ν-semifield† $\overline{F(\Lambda)}/o$ and $\sum_{i=1}^t \phi_o(|h_i|) \cap F = \{1\}$, then $h_1, ..., h_t$ are F-convexly dependent in $\overline{F(\Lambda)}$.*

Proof. Note that $\sum_{i=1}^t \phi_o(|h_i|) \cap F = \{1\}$ if and only if $\bigcap_{i=1}^t 1\text{-loc}(h_1) \cap 1\text{-loc}(o) \neq \emptyset$. Translating the variables by a point $a \in \bigcap_{i=1}^t 1\text{-loc}(h_1) \cap 1\text{-loc}(o)$, we may assume that the constant coefficient of each $\mathscr{L}$-monomial h_i is 1. Assume that $\phi_o(h_1), ..., \phi_o(h_t)$ are F-convexly dependent. We may assume that $\phi_o(h_1)$ is F-convexly dependent on $\phi_o(h_2), ..., \phi_o(h_t)$. This means by Definition 11.3.1 that we can take $h_{t+1} \in F$ for which $\phi_o(h_1) \in \langle \phi_o(h_2), ..., \phi_o(h_t), \phi_o(h_{t+1})\rangle$. Taking the pre-images of the quotient map yields

$$\langle h_1\rangle \cdot o \subseteq \langle h_2, ..., h_t, h_{t+1}\rangle \cdot o.$$

Take an $\mathscr{L}$-monomial g such that $1 + g$ generates o. By Corollary 4.1.21, there exists some $k \in \mathbb{N}$ such that

$$|h_1| + |1 + g| \leq_\nu \left(|h_2| + \cdots + |h_{t+1}| + |1 + g|\right)^k = |h_2|^k + \cdots + |h_{t+1}|^k + |1 + g|^k. \tag{11.6}$$

As $1 + g \geq_\nu 1$ we have that $|1 + g| \cong_\nu 1 + g$, and the right hand side of Equation (11.6) equals

$$|h_2|^k + \cdots + |h_{t+1}|^k + (1 + g)^k \cong_\nu |h_2|^k + \cdots + |h_{t+1}|^k + 1 + g^k \cong_\nu |h_2|^k + \cdots + |h_{t+1}|^k + g^k.$$

The last equality is due to the fact that $\sum |h_i|^k \geq_\nu 1$ so that 1 is absorbed. The same argument, applied to the left hand side of Equation (11.6), yields that

$$|h_1| + g \leq_\nu |h_2|^k + \cdots + |h_{t+1}|^k + g^k. \tag{11.7}$$

Assume on the contrary that h_1 is F-convexly independent of $\{h_2, ..., h_t\}$. Then

$$\langle h_1\rangle \not\subseteq \langle h_2, ..., h_{t+1}\rangle \cong_\nu \left\langle \sum_{i=2}^{t+1} |h_i| \right\rangle.$$

Thus for any $m \in \mathbb{N}$ there exists some $\mathbf{a}_m \in F^{(n)}$ such that

$$|h_1(\mathbf{a}_m)| >_\nu \left|\sum_{i=2}^{t+1} |h_i(\mathbf{a}_m)|\right|^m \cong_\nu \sum_{i=2}^{t+1} |h_i(\mathbf{a}_m)|^m.$$

Thus by equation (11.7) and the last observation we get that

$$\sum_{i=2}^{t} |h_i(\mathbf{a}_m)|^m + g(\mathbf{a}_m) <_\nu |h_1(\mathbf{a}_m)| + g(\mathbf{a}_m) \leq_\nu \sum_{i=2}^{t} |h_i(\mathbf{a}_m)|^k + g(\mathbf{a}_m)^k,$$

i.e., there exists some fixed $k \in \mathbb{N}$ such that for any $m \in \mathbb{N}$,

$$\sum_{i=2}^{t} |h_i(\mathbf{a}_m)|^m <_\nu \sum_{i=2}^{t} |h_i(\mathbf{a}_m)|^k + g(\mathbf{a}_m)^k. \tag{11.8}$$

For $m > k$, since $|\gamma|^k \leq_\nu |\gamma|^m$ for any $\gamma \in F$, we get that $\sum_{i=2}^t |h_i(\mathbf{a}_m)|^m \geq_\nu \sum_{i=2}^t |h_i(\mathbf{a}_m)|^k$. Write

$$g^k = g(1)g'.$$

Since g^k is an HP-kernel, $g(1)$ is the constant coefficient of g and g' is a Laurent monomial with coefficient 1.

According to the way the $\mathbf{a}_m$ were chosen, $\sum_{i=2}^t |h_i(\mathbf{a}_m)| > 1$ and $\sum_{i=2}^t |h_i(\mathbf{a}_m)|^m <_\nu g(\mathbf{a}_m)^k$, and thus $g(1) <_\nu g'(\mathbf{a}_{m_0})$ for large enough m_0. But $g'(\mathbf{a}_m^{-1}) \cong_\nu g'(\mathbf{a}_m)^{-1}$ so

$$g^k(\mathbf{a}_{m_0}^{-1}) \cong_\nu g(1)g'(\mathbf{a}_{m_0}^{-1}) \cong_\nu g(1)g'(\mathbf{a}_{m_0})^{-1} <_\nu 1.$$

Thus (11.8) yields $\sum_{i=2}^t |h_i(\mathbf{a}_{m_0}^{-1})|^m <_\nu \sum_{i=2}^t |h_i(\mathbf{a}_{m_0}^{-1})|^k$, a contradiction. □

Proposition 11.4.4. *Let R be a region ν-kernel of $\overline{F(\Lambda)}$. Let $\{h_1, ..., h_t\}$ be a set of $\mathscr{L}$-monomials such that $(R \cdot \langle h_1, ..., h_t\rangle) \cap F = \{1\}$. Then $h_1 \cdot R, ..., h_t \cdot R$ are F-convexly dependent in the quotient ν-semifield† $\overline{F(\Lambda)}/R$ if and only if $h_1, ..., h_t$ are F-convexly dependent in $\overline{F(\Lambda)}$.*

Proof. The 'if' part of the assertion follows from Proposition 11.4.2. Since $R = \prod_{i=1}^m o_i$ for suitable order ν-kernels $\{o_i\}_{i=1}^m$, the 'only if' part follows from Lemma 11.4.3 applied repeatedly to each of these o_i's. □

Proposition 11.4.5. *Let $R \in \mathcal{P}(\overline{F(\Lambda)})$ be a region ν-kernel. Then, for any set L of HS-fractions, we have*

$$\mathrm{d}_{\mathrm{conv}}(L) = \mathrm{d}_{\mathrm{conv}}(L \cdot R),$$

the right side taken in $\overline{F(\Lambda)}/R$.

Proof. $\mathrm{d}_{\mathrm{conv}}(L) \leq \mathrm{d}_{\mathrm{conv}}(R \cdot L)$ since $L \subseteq R \cdot L$. For the reverse inequality, let $\phi_R : \overline{F(\Lambda)} \to \overline{F(\Lambda)}/R$ be the quotient map. Since L is a sub-ν-semifield† of $\overline{F(\Lambda)}$, $\phi_R^{-1}(\phi_R(L)) = R \cdot \langle L\rangle$ by Theorem 4.1.12, and $\mathrm{d}_{\mathrm{conv}}(L) \geq \mathrm{d}_{\mathrm{conv}}(\phi_R(L))$ by Proposition 11.4.2, while $\mathrm{d}_{\mathrm{conv}}(\phi_R(L)) \geq \mathrm{d}_{\mathrm{conv}}(\phi_R^{-1}(\phi_R(L)))$ by Lemma 11.4.3. Thus $\mathrm{d}_{\mathrm{conv}}(L) \geq \mathrm{d}_{\mathrm{conv}}(R \cdot L)$. □

In this way we see that $K \mapsto \Omega(K)$ yields a homomorphism of ν-kernels. Hence Ω is a natural map in the sense of Definition 12.0.2, and we can apply Theorem 12.0.3.

Remark 11.4.6. Let R be a region ν-kernel and let

$$A = \{\langle g\rangle : \langle g\rangle \cdot \langle F\rangle \supseteq R \cdot \langle F\rangle\}.$$

Then $\mathrm{d}_{\mathrm{conv}}(\overline{F(\Lambda)}/R) = \mathrm{d}_{\mathrm{conv}}(A)$, in view of Remark 11.1.10 and Proposition 11.4.4. As $L_a \in A$ for any $a \in 1\text{-loc}(R) \neq \emptyset$ and $\mathrm{d}_{\mathrm{conv}}(L_a) = \mathrm{d}_{\mathrm{conv}}(\overline{F(\Lambda)})$, we conclude that $\mathrm{d}_{\mathrm{conv}}(A) = \mathrm{d}_{\mathrm{conv}}(\overline{F(\Lambda)})$.

We are ready to prove catenarity for $\mathrm{d}_{\mathrm{conv}}$.

Theorem 11.4.7. *If R is a region ν-kernel and L is an HS-kernel of $\overline{F(\Lambda)}$, then*

$$\mathrm{d}_{\mathrm{conv}}(\overline{F(\Lambda)}/LR) = \mathrm{d}_{\mathrm{conv}}(\overline{F(\Lambda)}) - \mathrm{d}_{\mathrm{conv}}(L).$$

In particular,

$$\mathrm{d}_{\mathrm{conv}}(\overline{F(\Lambda)}/LR) = n - \mathrm{d}_{\mathrm{conv}}(L).$$

Proof. $\overline{F(\Lambda)}/LR \cong (\overline{F(\Lambda)}/R)/(L \cdot R/R)$, by the third isomorphism theorem. Choose a basis for $\mathrm{HP}(\overline{F(\Lambda)}/R)$ containing a basis for $\mathrm{HP}(L \cdot R/R)$. Then $\mathrm{d}_{\mathrm{conv}}(L \cdot R/R) = \mathrm{d}_{\mathrm{conv}}(\phi_R(L))$, by Remark 11.4.6. But $L \cdot R \cap F = \{1\}$. Hence, by Proposition 11.4.4, $\mathrm{d}_{\mathrm{conv}}(\phi_R(L)) = \mathrm{d}_{\mathrm{conv}}(L)$. So

$$\mathrm{d}_{\mathrm{conv}}(\overline{F(\Lambda)}/LR) = \mathrm{d}_{\mathrm{conv}}(A) - \mathrm{d}_{\mathrm{conv}}(L) = \mathrm{d}_{\mathrm{conv}}(\overline{F(\Lambda)}) - \mathrm{d}_{\mathrm{conv}}(L).$$

Thus,

$$\mathrm{d}_{\mathrm{conv}}(\overline{F(\Lambda)}/LR) = \mathrm{d}_{\mathrm{conv}}(A) - \mathrm{d}_{\mathrm{conv}}(L) = n - \mathrm{d}_{\mathrm{conv}}(L).$$

□

11.5. **Recapitulation for dimensions.**

In conclusion, for every principal regular ν-kernel $\langle f\rangle \in P(\overline{F(\Lambda)})$, we have obtained explicit region ν-kernels

$$\{R_{1,1}, ..., R_{1,s}, R_{2,1}, ..., R_{2,t}\}$$

having trivial intersection, such that

$$\langle f\rangle = \bigcap_{i=1}^{s} K_i \cap \bigcap_{j=1}^{t} N_j$$

where $K_i = L_i \cdot R_{1,i}$ for $i = 1, ..., s$ and appropriate HS-kernels L_i and $N_j = B_j \cdot R_{2,j}$ for $j = 1, ..., t$ and appropriate bounded from below ν-kernels B_j. If $\langle f\rangle \in \mathcal{P}(\langle F\rangle)$, then, in view of Theorem 9.4.8 we can take $B_j = \langle F\rangle$ for every $j = 1, ..., t$. Note that over the various regions in $F^{(n)}$ corresponding to the region ν-kernels $R_{i,j}$, f is locally represented by distinct HS-fractions in $\mathrm{HSpec}(\overline{F(\Lambda)})$. In fact each region is defined so that the local HS-representation of f is given over the entire region. Thus the $R_{i,j}$'s defining the partition of the space can be obtained as a minimal set of regions over each of which $\langle f\rangle$ takes the form of an HS-kernel.

For each $j = 1, ..., t$, $\mathrm{d_{conv}}(N_j) = \mathrm{Hdim}(N_j) = 0$, since N_j contains no elements of $\mathrm{HP}(\overline{F(\Lambda)})$, implying

$$\mathrm{d_{conv}}(\overline{F(\Lambda)}/N_j) = \mathrm{Hdim}(\overline{F(\Lambda)}/N_j) = n.$$

For each $i = 1, ..., s$, $\mathrm{d_{conv}}(K_i) = \mathrm{d_{conv}}(L_i) = \mathrm{Hdim}(L_i) \geq 1$, implying

$$\mathrm{d_{conv}}(\overline{F(\Lambda)}/K_i) = \mathrm{Hdim}(\overline{F(\Lambda)}/K_i) = n - \mathrm{Hdim}(L_i) < n.$$

Remark 11.5.1. In view of the discussion in [29, §9.2], each term $\overline{F(\Lambda)}/L_i$ corresponds to the linear subspace of $F^{(n)}$ (in logarithmic scale) defined by the linear constraints endowed on the quotient $\overline{F(\Lambda)}/L_i$ by the HS-kernel L_i. One can think of these terms as an algebraic description of the affine subspaces locally comprising 1-loc(f).

12. The Jordan-Hölder theorem

Our final goal is to obtain a Jordan-Hölder theorem for ν-kernels. But there are too many ν-kernels for a viable theory in general, as discussed in [2], and anyway, as pointed out in the introduction, the order ν-kernels need not alter the geometric dimension,so we want to ignore them in this section. If we limit our set of ν-kernels to a given sublattice, of HS-kernels, one can use the Schreier refinement theorem [32] to obtain a version of the Jordan–Hölder Theorem.

Definition 12.0.2. $\mathcal{L}(\mathcal{S})$ denotes the lattice of ν-kernels of a ν-semifield† $\mathcal{S}$.

Ψ is a **natural map** if for each ν-semifield† $\mathcal{S}$, there is a lattice homomorphism $\Psi_{\mathcal{S}} : \mathcal{L}(\mathcal{S}) \to \mathcal{L}(\mathcal{S})$ such that $K \mapsto \Psi_{\mathcal{S}}(K)$ is a homomorphism of ν-kernels. We write $\Psi(\mathcal{S})$ for $\Psi_{\mathcal{S}}(\mathcal{L}(S))$, and call the ν-kernels in $\Psi(\mathcal{S})$ Ψ-**kernels**. (We delete $\mathcal{S}$ when it is unambiguous.)

A $\Psi(\mathcal{S})$-**simple** ν-kernel is a minimal Ψ-kernel $\neq \{1\}$. A $\Psi(\mathcal{S})$-**composition series** $\mathcal{C}(K, L)$ in $\Psi(\mathcal{S})$ from a ν-kernel K to a ν-subkernel L is a chain

$$K = K_0 \supset K_1 \supset \dots K_t = L$$

in Ψ such that each factor is Ψ-simple.

By Theorem 4.1.10 $\mathcal{C}(K, L)$ is equivalent to the $\Psi(\mathcal{S}/L)$-composition series

$$K/L \supset K_1/L \supset \cdots \supset K_t/L = 0$$

of K/L.

Given a Ψ-kernel K, we define its **composition length** $\ell(K)$ to be the length of a Ψ-composition series for K (presuming K has one). By definition, $\{1\}$ is the only Ψ-kernel of composition length 0. A nonzero Ψ-kernel K is simple iff $\ell(K) = 1$. The next theorem is a standard lattice-theoretic result of Schreier and Zassenhaus, yielding the Schreier-Jordan-Hölder Theorem, cf. [33, Theorem 3.11].

Theorem 12.0.3. *Suppose K has a composition series*

$$K = K_0 \supset K_1 \supset \cdots \supset K_t = 0,$$

which we denote as $\mathcal{C}$. Then:

(i) Any arbitrary finite chain of ν-subkernels

$$K = N_0 \supset N_1 \supset \cdots \supset N_k \supset 0$$

(denoted as $\mathcal{D}$), can be refined to a composition series equivalent to $\mathcal{C}$. In particular, $k \leq t$.

(ii) Any two composition series of K are equivalent.

(iii) $\ell(K) = \ell(N) + \ell(K/N)$ for every ν-subkernel N of K. In particular, every ν-subkernel and every homomorphic image of a ν-kernel with composition series has a composition series.

To apply this, we need the following observation.

Proposition 12.0.4. *Let L be an HS-kernel in $\overline{F(\Lambda)}$ with $1\text{-loc}(L) \neq \emptyset$. Let $\{h_1, ..., h_t\}$ be a set of $\mathscr{L}$-monomials in $\mathrm{HSpec}(\overline{F(\Lambda)})$ such that $\mathrm{Conv}(h_1, ..., h_t) = \mathrm{Conv}(L)$ and let $L_i = \langle h_i \rangle$. Then the chain*

$$L = \prod_{i=1}^{u} L_i \supseteq \prod_{i=1}^{u-1} L_i \supseteq \cdots \supseteq L_1 \supseteq \langle 1 \rangle. \tag{12.1}$$

of HS-kernels is strictly descending if and only if $h_1, ..., h_u$ are F-convexly independent.

Proof. ($\Rightarrow$) If h_u is F-convexly dependent on $\{h_1, ..., h_{u-1}\}$, then $L_u = \langle h_u \rangle \subseteq \prod_{i=1}^{u-1} L_i \cdot \langle F \rangle$. Assume that $L_u = \langle h_u \rangle \not\subseteq \prod_{i=1}^{u-1} L_i$. Then $\langle F \rangle \subseteq \prod_{i=1}^{u} L_i$, implying that $\prod_{i=1}^{u} L_i$ is not an HS-kernel. Thus $L_u = \langle h_u \rangle \subseteq \prod_{i=1}^{u-1} L_i$, and the chain is not strictly descending.

($\Leftarrow$) $1\text{-loc}(\prod_{i=1}^{t} L_i) \subseteq 1\text{-loc}(L) \neq \emptyset$ for every $0 \leq t \leq u$, implying that $(\prod_{i=1}^{t} L_i) \cap F = \{1\}$ for every $0 \leq t \leq u$ (for otherwise $1\text{-loc}(\prod_{i=1}^{t} L_i) = \emptyset$). If $\{h_1, ..., h_u\}$ is F-convexly independent then $L_u = \langle h_u \rangle \not\subseteq \prod_{i=1}^{u-1} L_i \cdot \langle F \rangle$. By induction, the chain (12.1) is strictly descending.

□

Theorem 12.0.5. *If $L \in \mathrm{HSpec}(\overline{F(\Lambda)})$, then $\mathrm{hgt}(L) = \mathrm{d}_{\mathrm{conv}}(L)$, cf. Definition 11.1.9. Moreover, every factor of a descending chain of HS-kernels of maximal length is an HP-kernel.*

Proof. By Proposition 12.0.4, the maximal length of a chain of HS-kernels descending from an HS-kernel L equals the number of elements in a basis of $\mathrm{Conv}(L)$; thus the chain is of unique length $\mathrm{d}_{\mathrm{conv}}(L)$, i.e., $\mathrm{hgt}(L) = \mathrm{d}_{\mathrm{conv}}(L)$. Moreover, by Theorem 4.1.12(2),

$$\prod_{i=1}^{j} L_i / \prod_{i=1}^{j-1} L_i \cong L_j / \left(L_j \cap \prod_{i=1}^{j-1} L_i \right).$$

Furthermore

$$\left(L_j \cdot (L_j \cap (\prod_{i=1}^{j-1} L_i)) \right) \cap F = \{1\},$$

since $L_j \cdot (L_j \cap \prod_{i=1}^{j-1} L_i) = L_j \cap \prod_{i=1}^{j} L_i \subset \prod_{i=1}^{j} L_i$ and $(\prod_{i=1}^{j} L_i) \cap F = \{1\}$. So the image of the HP-kernel L_j in $\overline{F(\Lambda)}/(L_j \cap (\prod_{i=1}^{j-1} L_i))$ is an HP-kernel. Thus, every factor of the chain is an HP-kernel. □

Corollary 12.0.6. $\mathrm{Hdim}(\overline{F(\Lambda)}) = \mathrm{d}_{\mathrm{conv}}(\overline{F(\Lambda)}) = n$.

Remark 12.0.7. If $R_1 \cap R_2 \cap \cdots \cap R_t = \{1\}$, then $\overline{F(\Lambda)}$ is a subdirect product

$$\overline{F(\Lambda)} = \overline{F(\Lambda)}/(R_1 \cap R_2 \cap \cdots \cap R_t) \hookrightarrow \prod_{i=1}^{t} \overline{F(\Lambda)}/R_i.$$

Then for any ν-kernel K of $\overline{F(\Lambda)}$, $R_1 \cap R_2 \cap \cdots \cap R_t \cap K = \bigcap_{i=1}^{t} (R_i \cap K) = \{1\}$ and, since K itself is an idempotent ν-semifield[†],

$$K = K/\bigcap_{i=1}^{t} (R_i \cap K) \cong \prod_{i=1}^{t} K/(R_i \cap K) \cong \prod_{i=1}^{t} R_i K/R_i.$$

References

1. A. Bertram and R. Easton, *The tropical nullstellensatz for congruences*, preprint (2013).
2. G. Birkhoff, *Transfinite subgroup series*, Bulletin of the American Mathematical Society 40 (12): 847850 (1934).
3. N. Bourbaki, *Commutative Algebra.* Paris and Reading, 1972.
4. D. Castella, *Algèbres de polynòmes tropicaux*, Annales mathématiques Blaise Pascal (2013).
5. P. Cohn, *Universal algebra*, Reidel, 1981.
6. G. Cortinas, C. Haesemeyer, M. Walker, and C. Weibel, *Toric varieties, monoid schemes, and cdh-descent*, (1970).
7. L.W. Cohen, C. Goffman, The topology of ordered Abelian groups. Trans. Amer. Math. Soc. (1949), 310–319.
8. S. Gaubert, M. Akian, and A. Guterman, *Linear independence over tropical semirings and beyond*, Tropical and Idempotent Mathematics, G.L. Litvinov and S.N. Sergeev, (eds.) **495** (2009), 1–38.
9. J. Giansiracusa and N. Giansiracusa, *Equations of tropical varieties*, arXiv:1308.0042v1 [math.AG] (2013).
10. J. Golan, *The theory of semirings with applications in mathematics and theoretical computer science*, vol. 54, Pitman Monographs and Surveys in Pure and Applied Mathematics, Longman Scientific and Technical, 1992.
11. C. Holland, *Extensions of ordered groups and sequence completion*, Trans. Amer. Math. Soc. 107 (1963), 71–82.
12. H. Hutchins and H. Weinert, *Homomorphisms and kernels of semifields*, Periodica Mathematica Hungarica **21 (2)** (1990), 113–152.
13. Z. Izhakian, M. Knebusch, and L. Rowen, *Layered tropical mathematics*, J. Algebra, to appear, preprint (2013), available at `http://arxiv.org/pdf/0912.1398.pdf`.
14. ———, *Supertropical semirings and supervaluations*, J. Pure Appl. Algebra **215(10)** (2011), 2431–2463.
15. ———, *Categorical notions of layered tropical algebra and geometry*, Algebraic and combinatorial aspects of tropical geometry, 191234, Contemp. Math., 589, Amer. Math. Soc., Providence, RI, 2013.

16. ———, *Algebraic structures of tropical mathematics, in Tropical and Idempotent Mathematics and Applications*, G.L. Litvinov and S.N.Sergeev, editors, Cont. Math. 616 arXiv:1305.3906v1 [math.RA] 16 May 2013
17. Z. Izhakian and L. Rowen, *Supertropical algebra*, Adv. in Math **225** (2010), 2222–2286.
18. ———, *Supertropical matrix algebra*, Israel J. Math. **182** (2011), 383–424.
19. ———, *Ideals of polynomial semirings in tropical mathematics*, J. Algebra Appl. **12(2)** (2013), 125–143.
20. ———, *Coordinate semirings of tropical varieties*, preprint (2014).
21. N. Jacobson, *Basic Algebra II*, Freeman, 1980.
22. D. Joo and K. Mincheva, *Prime congruences of idempotent semirings and a Nullstellensatz for tropical polynomials*, arXiv:1408.3817v1 [math.AC] (17 Aug 2014).
23. V.M. Kopytov, Lattice ordered groups (Russian), Nauka, Moscow (1984)
24. D. Maclagan and B. Sturmfels, *Introduction to Tropical Geometry*, To be published by the American Mathematical Society.
25. D. Maclagan and F. Rinćon, *Tropical schemes, tropical cycles, and valuated matroids*, arXiv:1401.4654v1 [math.AG] (2014).
26. A. Macpherson, *Skeleta in non-archimedean and tropical geometry*, arXiv:1311.0502v1 [math.AG] (2013).
27. D. Marker, Model theory : An introduction Springer Graduate texts in mathematics; 217, 2002.
28. T. Perri, *Algebraic aspects in tropical mathematics*, preprint, available at `http://arxiv.org/pdf/1305.2764v2.pdf` (2013).
29. T. Perri and L. Rowen, *Kernels in tropical geometry*, preprint, `http://arxiv.org/pdf/14v2.pdf` (2013).
30. T. Perri and L. Rowen, *A tropical Krull-Schmidt theorem* preprint, available at `http://arxiv.org/pdf/1408.4757` (2014).
31. B. Plotkin, E. Aladova, and E. Plotkin, *Algebraic logic and logically-geometric types in varieties of algebras*, Journal of Algebra and its Applications **12 (2)** (2013), 23pp.
32. J. Rotman, *An introduction to the theory of groups*, New York: Springer-Verlag. ISBN 0-387-94285-8.Rotman, Joseph (1994).
33. L. Rowen, *Graduate Algebra: Commutative View*, AMS Graduate Studies in Mathematics **73** 2006.
34. D. Marker, Model theory : An introduction Springer Graduate texts in mathematics; 217, 2002.
35. G.E. Sacks, Saturated Model Theory, Mathematical Lecture Noets 80 Benjamin, 1972.
36. S. Steinberg, *Lattice-ordered Rings and Modules*, Springer Science, 2010.
37. E. Vechtomov and A. Cheraneva, *Semifields and their Properties*, Journal of Mathematical Sciences **163** (2009), 625–661.
38. E. Vechtomov and D. Chuprakov, *The principal kernels of semifields of continuous positive functions*, Journal of Mathematical Sciences **163** (2009), 500–514.
39. H. Weinert and R. Wiegandt, *On the structure of semifields and lattice-ordered groups*, Periodica Mathematics Hungarica **32(1-2)** (1996), 129–147.
40. ———, *A new Kurosh-Amitsur radical theory for proper semifields*, Mathematica Pannonica **14(1)** (2003), 3–28.
41. T.Y. Yu, Balancing conditions in global tropical geometry, arXiv:1304.2251v1 [math.AG], 2013

E-mail address: `talperri@zahav.net.il`
E-mail address: `rowen@math.biu.ac.il`